arXiv:1512.01500v4 [math.DS] 6 Jun 2016

The geometry of model spaces for probability-preserving actions of sofic groups

Tim Austin

Abstract

Bowen's notion of sofic entropy is a powerful invariant for classifying probability-preserving actions of sofic groups. It can be defined in terms of the covering numbers of certain metric spaces associated to such an action, the 'model spaces'.

The metric geometry of these model spaces can exhibit various interesting features, some of which provide other invariants of the action. This paper explores an approximate connectedness property of the model spaces, and uses it give a new proof that certain groups admit factors of Bernoulli shifts which are not Bernoulli. This was originally proved by Popa. Our proof covers fewer examples than his, but provides additional information about this phenomenon.

MSC(2010): 37A35 (primary), 28D20, 37A25, 37A50, 60J10 (secondary)

Kewords: Sofic entropy, Bernoulli system, Popa factor, model spaces, connectedness

Contents

1 Introduction **2**

2 Connected model spaces **6**

3 Model-surjective factor maps and Theorem A **13**

4 Connected model spaces for inverse limits **21**

5 Connected model spaces for Bernoulli systems **22**

6 Analysis of Popa factors **28**

1 Introduction

Let G be a discrete sofic group, (X, μ) a standard probability space, and $T : G \curvearrowright X$ a measurable action preserving μ. The triple (X, μ, T) is called a G-**system** or just a **system**. Fix a choice of sofic approximation $\Sigma = (\sigma_n)_{n\geqslant 1}$ to G, where each σ_n is a map from G to $\mathrm{Sym}(V_n)$ for some finite set V_n.

The sofic entropy of (X, μ, T) relative to this sofic approximation gives a generalization of the classical Kolmogorov-Sinai entropy which is still available if G is not amenable. If (X, μ, T) has a finite generating partition, then its sofic entropy rel Σ was defined by Lewis Bowen in [3], and this definition was generalized to any G-system by Kerr and Li in [13, 14]. In case G is amenable, the agreement between sofic entropy rel Σ and Kolmogorov-Sinai entropy was shown in [6].

The present paper introduces a new invariant of G-systems, constructed from similar ingredients to the sofic entropy. Our starting point is the formulation of sofic entropy given in [1], which slightly modifies [14] without changing the resulting entropy values.

In the first place, we consider a special class of G-systems. A **metric** G-**process** is a quadruple $(\mathcal{X}^G, \mu, S, d)$ in which $\mathcal{X}$ is another standard measurable space, d is a compact metric on $\mathcal{X}$ whose Borel sets are the σ-algebra of $\mathcal{X}$, S is the right-shift action on $\mathcal{X}^G$, and μ is an S-invariant probability. Then $\mathrm{h}_\Sigma(\mu)$ is defined in terms of the covering numbers of certain metric spaces of 'good models' for μ over Σ. The 'good models' used in [1] differ slightly from their earlier counterparts in [13, 14], but it is shown in [1, Subsection 3.2] that one always obtains the same values for sofic entropy as in those earlier papers. The resulting entropy value $\mathrm{h}_\Sigma(\mu)$ does not depend on the choice of d and is an isomorphism-invariant of the G-system $(\mathcal{X}^G, \mu, S)$. Since any G-system is isomorphic to a shift-system, this lets us extend the definition of h_Σ unambiguously to arbitrary G-systems.

In the approach of [1], and similarly to the other approaches, 'good models' for $(\mathcal{X}^G, \mu, S, d)$ are certain elements of the finite-dimensional Cartesian powers $\mathcal{X}^{V_n}$. Each $\mathbf{x} \in \mathcal{X}^{V_n}$ has an 'empirical distribution' $P^{\sigma_n}_{\mathbf{x}}$ which may be used to describe its 'local statistics' over certain subsets of V_n: see Subsection 2.1. It is a 'good model' if its empirical distribution is close to μ in the weak* topology on $\mathrm{Prob}(\mathcal{X}^G)$. For any weak*-neighbourhood $\mathcal{O}$ of μ, let

$$\Omega(\mathcal{O}, \sigma_n) := \{\mathbf{x} \in \mathcal{X}^{V_n} : \ P^{\sigma_n}_{\mathbf{x}} \in \mathcal{O}\}$$

be the set of models which are 'good' according to this neighbourhood $\mathcal{O}$.

Also, for each n, let $d^{(V_n)}$ be the Hamming average of copies of d on $\mathcal{X}^{V_n}$:

$$d^{(V_n)}(\mathbf{x}, \mathbf{x}') = \frac{1}{|V_n|} \sum_{v \in V_n} d(x_v, x'_v) \quad \text{for } \mathbf{x} = (x_v)_{v \in V_n},\ \mathbf{x}' = (x'_v)_{v \in V_n}.$$

Sofic entropy is defined in terms of certain asymptotic geometric features of the sequences of metric spaces

$$\big(\Omega(\mathcal{O}, \sigma_n), d^{(V_n)}\big).$$

Specifically, it is the quantity

$$\mathrm{h}_\Sigma(\mu) := \sup_{\delta > 0} \inf_{\mathcal{O}} \limsup_{n \longrightarrow \infty} \frac{1}{|V_n|} \log \mathrm{cov}_\delta\big(\Omega(\mathcal{O}, \sigma_n), d^{(V_n)}\big),$$

where cov_δ is the δ-covering number of the given metric space and $\mathcal{O}$ runs over all weak*-neighbourhoods of μ.

The original motivation for studying sofic entropy was the classification of Bernoulli systems: see [3, 12]. It has quickly found numerous other applications, and suggests many directions for further research. One of these is to look for other properties of the metric spaces $(\Omega(\mathcal{O}, \sigma_n), d^{(V_n)})$ that carry some dynamical information about μ, and might give rise to other invariants.

The present paper focuses on one such property, which we refer to as 'connected model spaces rel Σ'. It is an approximate kind of connectedness for the spaces $\Omega(\mathcal{O}, \sigma_n)$. In case $\mathcal{X}$ is finite, these metric spaces are discrete, so we do not ask for their connectedness as in classical point-set topology. Rather, we consider whether connectedness holds in a suitable asymptotic sense as $\mathcal{O}$ becomes smaller and $n \longrightarrow \infty$. The precise notion we need is given in Definition 2.5.

Once that definition has been made, we can show that the property of connected model spaces rel Σ is an isomorphism-invariant. In fact we prove more than this. Some factor maps do not preserve the property of connected model spaces, but we introduce a class of factor maps which do, and show that they include all isomorphisms.

Given a factor map

$$\Phi : (\mathcal{X}^G, \mu, S) \longrightarrow (\mathcal{Y}^G, \nu, S),$$

and also compact metrics $d_\mathcal{X}$ on $\mathcal{X}$ and $d_\mathcal{Y}$ on $\mathcal{Y}$ which generate their σ-algebras, the developments in [1] include the construction of suitable 'approximating maps'

from $\mathcal{X}^{V_n}$ to $\mathcal{Y}^{V_n}$. These send good models for μ to good models for ν, up to various errors that must be carefully controlled in terms of the metrics $d_{\mathcal{X}}^{(V_n)}$ and $d_{\mathcal{Y}}^{(V_n)}$. Put roughly, we say that Φ is 'model-surjective rel Σ' if *all* good models for $(\mathcal{Y}^G, \nu, S)$ are close to images of good models for $(\mathcal{X}^G, \mu, S)$ under these approximating maps. This is made precise in Definition 3.1.

Having introduced this special property of factor maps, we will show that it is preserved by composition, and that all isomorphisms have this property. Then we prove the following.

Theorem A *Let Φ be a factor map as above. If $(\mathcal{X}^G, \mu, S, d_{\mathcal{X}})$ has connected model spaces rel Σ and Φ is model-surjective rel Σ, then $(\mathcal{Y}^G, \nu, S, d_{\mathcal{Y}})$ also has connected model spaces rel Σ.*

In particular, having connected model spaces rel Σ is a property only of the shift-system $(\mathcal{X}^G, \mu, S)$, not depending on the choice of $d_{\mathcal{X}}$, and is an isomorphism-invariant.

Having proved Theorem A, the property of having connected model spaces rel Σ can be extended unambiguously to any G-system, by picking an isomorphism from it to G-process.

Although factor maps need not preserve the connectedness of model spaces rel Σ, it turns out inverse limits do.

Theorem B *If all members of an inverse sequence of G-systems have connected model spaces rel Σ, then so does the inverse limit.*

The rest of the paper is given to some examples of connected and non-connected model spaces. In the first place, we have the following.

Theorem C *For any sofic group G and sofic approximation Σ, Bernoulli systems over G have connected model spaces rel Σ.*

On the other hand, for some groups G and sofic approximations Σ there are factors of Bernoulli shifts that do not have connected model spaces (and so, in particular, not all factor maps are model-surjective). The last section of the paper is given to a family of such examples.

Let m be the Haar probability measure on the circle $\mathbb{T}$, and consider the factor map of G-systems

$$(\mathbb{T}^G, m^{\times G}, S) \longrightarrow (X, \mu, T) \tag{1}$$

defined by forming the quotient by the diagonal subgroup of $\mathbb{T}^G$: that is,

$$X = \mathbb{T}^G / \{(\ldots, \theta, \theta, \ldots) : \theta \in \mathbb{T}\},$$

μ is the Haar measure on this quotient group, and T is the quotient action.

Theorem D *If G is a residually finite group with Kazhdan's property (T), then G has a sofic approximation Σ relative to which the factor system (X, μ, T) constructed above does not have connected model spaces.*

In particular, this implies that (X, μ, T) is not Bernoulli.

These examples form a special case of a construction of non-Bernoulli factors of Bernoulli shifts due to Popa [18], so for general G we refer to (X, μ, T) as the **Popa factor** of $(\mathbb{T}^G, m^{\times G}, S)$. These factors show a striking difference between ergodic theory for amenable and non-amenable groups: among the former, factors of Bernoulli shifts are always still Bernoulli [17]. Popa's proof in [18] gives examples for a considerably larger range of groups G than those in Theorem D. It rests on calculations of the cohomology of these actions, generalizing work of Popa and Sasyk [19]. Our proof of Theorem D also uses some (more elementary) cohomology theory, and it covers fewer examples, but it gives a geometric interpretation to the non-Bernoullicity of these factors.

As pointed out to me by Brandon Seward, combining Theorems B, C and D immediately gives the following.

Corollary D′ *If G is as in Theorem D, then there are factors of Bernoulli G-systems which are not inverse limits of Bernoulli G-systems.* □

This contrasts with actions of amenable groups, among which inverse limits of Bernoulli systems, like factors, are also always still Bernoulli [17].

In the reverse direction, there are many groups which admit inverse limits of Bernoulli shifts which are not factors of Bernoulli shifts. Indeed, if G has a non-amenable free subgroup, then Bowen has shown [4] that every Bernoulli shift over G factors onto every other Bernoulli shift (this property may in fact hold for all non-amenable G: see [5, Corollary 1.6] for partial progress). One may therefore form an inverse sequence

$$\cdots \longrightarrow (A_3^G, p_3^{\times G}, S) \longrightarrow (A_2^G, p_2^{\times G}, S) \longrightarrow (A_1^G, p_1^{\times G}, S)$$

of Bernoulli systems in which each A_n is a finite alphabet and the Shannon entropies $\mathrm{H}(p_n)$ have finite sum. Now the inverse limit of this sequence has generating partitions of arbitrarily small Shannon entropy, by joining the generating

partitions of all systems that are sufficiently high up the sequence. As a result, that inverse limit has Rokhlin entropy zero, and hence also sofic entropy at most zero [21, 22]. On the other hand, every factor of a Bernoulli shift has positive sofic entropy [11].

Our tools can also be used to prove another strengthening of Theorem D. Recall that a factor map $\Phi : (X, \mu, T) \longrightarrow (Y, \nu, S)$ of G-systems is **complemented** if there is another factor map $\Psi : (X, \mu, T) \longrightarrow (Z, \theta, R)$ such that the combined map (Φ, Ψ) is a measure-theoretic isomorphism

$$(X, \mu, T) \xrightarrow{\cong} (Y \times Z, \nu \times \theta, S \times R).$$

In this case (Y, ν, S) is a **complemented factor** of (X, μ, T).

The Popa factors give examples of factor maps of Bernoulli shifts that are not model-surjective. However, it turns out that all complemented factor maps of Bernoulli shifts (and actually of a much more general class of systems) are model-surjective: see Theorem 6.8. Combined with Theorems A and C, this shows the following.

Corollary D″ *If G is as in Theorem D then the Popa factor (X, μ, T) is not a complemented factor of any Bernoulli shift.*

This corollary was also suggested to me by Brandon Seward. It can also be proved using Popa's cohomological approach, as has been shown to me by Yongle Jiang. The main new ingredient is the fact that, if (X, μ, T) is a complemented factor of any Bernoulli shift $(\mathcal{X}^G, \nu^{\times G}, S)$, then the resulting inclusion homomorphism from the degree-one, $\mathbb{T}$-valued cohomology of the former into the latter is injective. Similar arguments can be found in Jiang's recent paper [9].

Finally, the above discussion suggests the following.

Question 1.1. *Is it true that for any group G, any complemented factor of a Bernoulli system is Bernoulli?* ◁

I do not know that the answer is Yes for any non-amenable group.

2 Connected model spaces

2.1 Model spaces and maps between them

This part of the paper follows closely the approach to sofic entropy developed in [1, Part I], and mostly uses the same notation.

We use 'big-O' and 'little-o' notation without further comment. Among real numbers, we sometimes write '$a \approx_\varepsilon b$' in place of '$|a-b| < \varepsilon$'.

If (X, d_X) and (Y, d_Y) are two metric spaces, $\varepsilon > 0$ and $L < \infty$, then a map $\varphi : X \longrightarrow Y$ is **ε-almost L-Lipschitz** if

$$d_Y(\varphi(x), \varphi(x')) \leqslant \varepsilon + L d_X(x, x') \quad \forall x, x' \in X.$$

A map is **ε-almost Lipschitz** if this holds for some L.

We next recall some of the basic definitions and results from [1, Part I]. They are mostly small modifications to [3, 14]. Suppose that $(\mathcal{X}^G, \mu, S, d)$ is a **metric G-process**, meaning that $(\mathcal{X}, d)$ is a compact metric space, S is the right-shift action of G on $\mathcal{X}^G$, and $\mu \in \mathrm{Prob}(\mathcal{X}^G)$ is S-invariant.

Given a finite set V and a map $\sigma : G \longrightarrow \mathrm{Sym}(V)$, and also $v \in V$ and $\mathbf{x} \in \mathcal{X}^V$, we define the **pullback name of $\mathbf{x}$ at** v by

$$\Pi_v^\sigma(\mathbf{x}) := (x_{\sigma^g(v)})_{g \in G} \in \mathcal{X}^G.$$

In terms of these, the **empirical distribution** of $\mathbf{x}$ is

$$P_{\mathbf{x}}^\sigma := \frac{1}{|V|} \sum_{v \in V} \delta_{\Pi_v^\sigma(\mathbf{x})}.$$

For any w*-neighbourhood $\mathcal{O}$ of μ in $\mathrm{Prob}(\mathcal{X}^G)$, the **$\mathcal{O}$-good models for** μ are the elements of

$$\Omega(\mathcal{O}, \sigma_n) := \{\mathbf{x} : \ P_{\mathbf{x}}^\sigma \in \mathcal{O}\}.$$

Motivation for these concepts and their use in defining sofic entropy can be found in [1, Subsection 3.1].

Given any map $\psi : \mathcal{X}^G \longrightarrow \mathcal{Y}$ and also a map $\sigma : G \longrightarrow \mathrm{Sym}(V)$ for some finite set V, we define the associated map $\psi^\sigma : \mathcal{X}^V \longrightarrow \mathcal{Y}^V$ by

$$\psi^\sigma(\mathbf{x}) := \big(\psi(\Pi_v^\sigma(\mathbf{x}))\big)_{v \in V},$$

as in [1, Subsection 4.2].

Now let $\Phi = \varphi^G : (\mathcal{X}^G, \mu, S, d_\mathcal{X}) \longrightarrow (\mathcal{Y}^G, \nu, S, d_\mathcal{Y})$ be a factor map of metric G-processes, where we use the same notation as in [1, Subsection 4.1]. As in that reference, an **η-almost Lipschitz** (or **η-AL**) **approximation to φ rel** $(\mu, d_\mathcal{X}, d_\mathcal{Y})$ is a measurable map $\psi : \mathcal{X}^G \longrightarrow \mathcal{Y}$ with the following properties.

i) The map ψ approximates φ in the sense that

$$\int d_\mathcal{Y}(\varphi(x), \psi(x))\, \mu(\mathrm{d}x) < \eta. \tag{2}$$

ii) There is a finite $D \subseteq G$ such that ψ is D-local: that is, it depends only on coordinates in D.

iii) There is a D-local open subset $U \subseteq \mathcal{X}^G$ such that $\mu(U) > 1 - \eta$ and such that $\psi|U$ is η-almost Lipschitz from $d_{\mathcal{X}}^{(D)}$ to $d_{\mathcal{Y}}$.

These always exist for every η [1, Lemma 4.3]. In the applications below, we will more often consider sequences of such approximations: an **almost Lipschitz approximating sequence for φ rel** $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$ is a sequence of maps $\psi_k : \mathcal{X}^G \longrightarrow \mathcal{Y}$ which are η_k-AL approximations to φ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$ for some positive parameters $\eta_k \longrightarrow 0$. This situation is denoted by $\psi_k \xrightarrow{\text{aL}} \varphi$.

Although we will be using such maps to study isomorphism-invariant properties of systems, the definitions above depend crucially on the choice of the compact metrics $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$.

Later in this section, we will need some estimates from [1] concerning AL approximations and the corresponding maps on model spaces. Suitable versions are recalled in the following lemmas.

Lemma 2.1 (Good models are mapped to good models: [1, Proposition 4.10])**.** *For every w*-neighbourhood $\mathcal{N}$ of ν there is an $\eta > 0$ with the following property. If ψ is an η-AL approximation to φ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$, then there is a w*-neighbourhood $\mathcal{O}$ of μ such that*

$$\psi^{\sigma_n}\big(\Omega(\mathcal{O}, \sigma_n)\big) \subseteq \Omega(\mathcal{N}, \sigma_n)$$

for all sufficiently large n. □

Lemma 2.2 (Almost Lipschitz action on model spaces: [1, Lemma 4.9])**.** *If ψ is an η-AL approximation to φ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$, then there are $K < \infty$ and a w*-neighbourhood $\mathcal{O}$ of μ such that*

$$\psi^{\sigma_n}|\Omega(\mathcal{O}, \sigma_n)$$

is (3η)-almost K-Lipschitz from $d_{\mathcal{X}}^{(V_n)}$ to $d_{\mathcal{Y}}^{(V_n)}$ for all sufficiently large n. □

For our later application of these results, it is convenient to combine them into a single assertion in terms of bases for the w*-topologies around μ and ν. Since those topologies are metrizable, there exist such bases which are decreasing sequences of open neighbourhoods. We will frequently use this fact without further explanation.

Corollary 2.3. *Suppose that $\psi_k \xrightarrow{\text{aL}} \varphi$ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$. Then there are*

- *parameters* $\varepsilon_k \downarrow 0$ *and* $K_k < \infty$*,*
- *a sequence of w*-neighbourhoods* $\mathcal{O}_1 \supseteq \mathcal{O}_2 \supseteq \dots$ *of* μ*,*
- *and a base of w*-neighbourhoods* $\mathcal{N}_1 \supseteq \mathcal{N}_2 \supseteq \dots$ *at* ν

such that for each k *both of the following hold for all sufficiently large* n*:*

$$\psi_k^{\sigma_n}\big(\Omega(\mathcal{O}_k, \sigma_n)\big) \subseteq \Omega(\mathcal{N}_k, \sigma_n)$$

and

$$\psi_k^{\sigma_n}|\Omega(\mathcal{O}_k, \sigma_n) \quad \textit{is } \varepsilon_k\textit{-almost } K_k\textit{-Lipschitz from } d_{\mathcal{X}}^{(V_n)} \textit{ to } d_{\mathcal{Y}}^{(V_n)}.$$

Proof. Let $\eta_k \downarrow 0$ be a sequence such that ψ_k is an η_k-AL approximation to φ rel μ for every k, and set $\varepsilon_k := 3\eta_k$ for each k.

Let $\mathcal{M}_1 \supseteq \mathcal{M}_2 \supseteq \dots$ be a base of w*-neighbourhoods at ν with $\mathcal{M}_1 = \mathrm{Prob}(\mathcal{Y}^G)$. By Lemma 2.1, for each $j \in \mathbb{N}$ there are $k_0(j) \in \mathbb{N}$ and w*-neighbourhoods $\mathcal{O}'_{j,k}$ for each $k \geqslant k_0(j)$ such that for every $k \geqslant k_0(j)$ we have

$$\psi_k^{\sigma_n}\big(\Omega(\mathcal{O}'_{j,k}, \sigma_n)\big) \subseteq \Omega(\mathcal{M}_j, \sigma_n) \quad \text{for all sufficiently large } n.$$

We may take $k_0(1) = 1$ because $\mathcal{M}_1 = \mathrm{Prob}(\mathcal{Y}^G)$, and we may also assume that $k_0(j) \longrightarrow \infty$ as $j \longrightarrow \infty$. Now let

$$\mathcal{N}_k := \mathcal{M}_{\max\{j\colon\, k_0(j) \leqslant k\}}$$

and choose w*-neighbourhoods $\mathcal{O}'_k$ of μ satisfying

$$\mathcal{O}'_k \subseteq \bigcap_{j\colon\, k_0(j) \leqslant k} \mathcal{O}'_{j,k}.$$

By forming running intersections we may also assume that $\mathcal{O}'_1 \supseteq \mathcal{O}'_2 \supseteq \dots$. For each k, letting j be the largest integer for which $k_0(j) \leqslant k$, it follows that

$$\psi_k^{\sigma_n}\big(\Omega(\mathcal{O}'_k, \sigma_n)\big) \subseteq \psi_k^{\sigma_n}\big(\Omega(\mathcal{O}'_{j,k}, \sigma_n)\big),$$

and this is contained in $\Omega(\mathcal{M}_j, \sigma_n) = \Omega(\mathcal{N}_k, \sigma_n)$ for all sufficiently large n.

On the other hand, Lemma 2.2 gives values $K_k < \infty$ and w*-neighbourhoods $\mathcal{O}''_k$ of μ such that $\psi_k^{\sigma_n}|\Omega(\mathcal{O}''_k, \sigma_n)$ is ε_k-almost K_k-Lipschitz for all sufficiently large n, for every $k \geqslant 1$. Once again we may assume that $\mathcal{O}''_1 \supseteq \mathcal{O}''_2 \supseteq \dots$.

Finally, let $\mathcal{O}_k := \mathcal{O}'_k \cap \mathcal{O}''_k$ for each k. □

The deduction of this corollary from Lemmas 2.1 and 2.2 is similar to the deduction of 'sequence versions' of those lemmas in [1, Subsection 4.3]. But it seems easier to make this deduction from scratch here, rather than adapting the corollaries in that subsection.

Corollary 2.3 has a further consequence that is worth recording by itself.

Corollary 2.4. *If* $\psi_k \xrightarrow{\text{aL}} \varphi$ *rel* $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$ *and if* $\mathbf{x}_n \in \mathcal{X}^{(V_n)}$ *is a sequence of models such that* $P^{\sigma_n}_{\mathbf{x}_n} \xrightarrow{\text{weak}^*} \mu$*, then*

$$P^{\sigma_n}_{\psi^{\sigma_n}_{k_n}(\mathbf{x}_n)} \xrightarrow{\text{weak}^*} \nu$$

whenever the sequence $k_1 \leqslant k_2 \leqslant \dots$ *grows sufficiently slowly.*

Proof. This really uses only the first conclusion of Corollary 2.3: there are sequences $\mathcal{O}_1 \supseteq \mathcal{O}_2 \supseteq \dots$ and $\mathcal{N}_1 \supseteq \mathcal{N}_2 \supseteq \dots$ as in that corollary such that for all k we have

$$\psi^{\sigma_n}_k\big(\Omega(\mathcal{O}_k, \sigma_n)\big) \subseteq \Omega(\mathcal{N}_k, \sigma_n) \quad \text{for all sufficiently large } n.$$

Since $P^{\sigma_n}_{\mathbf{x}_n} \xrightarrow{\text{weak}^*} \mu$, for every k there is an $N(k)$ such that $\mathbf{x}_n \in \Omega(\mathcal{O}_k, \sigma_n)$ for all $n \geqslant N(k)$. Provided the sequence $k_1 \leqslant k_2 \leqslant \dots$ grows sufficiently slowly, it follows that $\mathbf{x}_n \in \Omega(\mathcal{O}_{k_n}, \sigma_n)$ for all sufficiently large n, and also that

$$\psi^{\sigma_n}_{k_n}\big(\Omega(\mathcal{O}_{k_n}, \sigma_n)\big) \subseteq \Omega(\mathcal{N}_{k_n}, \sigma_n) \quad \text{for all sufficiently large } n.$$

These together imply that

$$\psi^{\sigma_n}_{k_n}(\mathbf{x}_n) \in \Omega(\mathcal{N}_{k_n}, \sigma_n) \quad \text{for all sufficiently large } n,$$

and so their empirical distributions converge to ν, since $\mathcal{N}_1 \supseteq \mathcal{N}_2 \supseteq \dots$ is a base for the w*-topology at ν. □

2.2 Connected model spaces

Let (Y, d_Y) be a metric space, let $x, y \in Y$, and let $\delta > 0$. A δ-**path from** x **to** y is a finite sequence

$$x = x_0, x_1, \dots, x_\ell = y$$

in Y such that $d_Y(x_i, x_{i+1}) < \delta$ for every $i \in \{0, \dots, \ell - 1\}$. The integer ℓ is the **length** of this δ-path. If $A \subseteq Y$ and $\delta > 0$, then A is δ-**connected** (**according to** d_Y) if for any $x, y \in A$ there is a δ-path from x to y contained in A.

We are now ready to define our new property of metric G-processes:

Definition 2.5 (Connected model spaces rel Σ)**.** *Let $(\mathcal{X}^G, \mu, S, d)$ be a metric G-process. It has* ***connected model spaces rel*** *Σ if the following holds:*

If $n_1 < n_2 < \ldots$, and $\mathbf{x}_i, \mathbf{y}_i \in \mathcal{X}^{V_{n_i}}$ are two sequences satisfying

$$P^{\sigma_{n_i}}_{\mathbf{x}_i}, P^{\sigma_{n_i}}_{\mathbf{y}_i} \overset{\text{weak}^*}{\longrightarrow} \mu,$$

then there are a sequence $\delta_i \downarrow 0$ and a sequence of δ_i-paths

$$\{\mathbf{x}_i = \mathbf{x}_{i,0},\ \mathbf{x}_{i,1},\ \ldots,\ \mathbf{x}_{i,\ell_i} = \mathbf{y}_i\} \subseteq \mathcal{X}^{V_{n_i}}$$

(according to the metrics $d^{(V_{n_i})}$) such that for any w-neighbourhood $\mathcal{O}$ of μ we have*

$$\{\mathbf{x}_{i,j} :\ 0 \leqslant j \leqslant \ell_i\} \subseteq \Omega(\mathcal{O}, \sigma_{n_i}) \tag{3}$$

for all sufficiently large i.

This definition is made more complicated by the allowance of an arbitrary subsequence $n_1 < n_2 < \ldots$. This is because of cases in which there are some other subsequence $n'_1 < n'_2 < \ldots$ and a w*-neighbourhood $\mathcal{O}$ such that $\Omega(\mathcal{O}, \sigma_{n'_i}) = \varnothing$ for all i. Such cases should still be called 'connected' if any two sufficiently good models, for a sufficiently large value of n, can be joined by a δ-path consisting of fairly good models. This requirement simply ignores any values n' for which $\sigma_{n'}$ admits no good models at all. However, if there is a subsequence $n'_1 < n'_2 < \ldots$ as above, then there can be no sequences $\mathbf{x}_n, \mathbf{y}_n \in \mathcal{X}^{V_n}$ defined for *all* integers n which satisfy $P^{\sigma_n}_{\mathbf{x}_n}, P^{\sigma_n}_{\mathbf{y}_n} \overset{\text{weak}^*}{\longrightarrow} \mu$, and we must pass to a subsequence which eventually avoids the n'_is. This will be made clearer by the proof of Proposition 2.6 below.

Definition 2.5 can be re-written more directly in terms of connectedness properties of the model spaces $\Omega(\mathcal{O}, \sigma_n)$. At first sight, it seems similar to requiring that these spaces are δ-connected for all sufficiently large n, but that impression is not quite correct. In order to fix it, we need a slightly more complicated notion.

If (Y, d_Y) is any metric space and $A \subseteq B \subseteq Y$ is a nested pair of subsets, then the pair (A, B) is **relatively δ-connected (according to d_Y)** if for any $x, y \in A$ there is a δ-path from x to y contained in B.

Proposition 2.6. *If $(\mathcal{X}^G, \mu, S, d)$ is a metric G-process, then the following are equivalent.*

1. *It has connected model spaces rel Σ.*

2. *For every $\delta > 0$ and every w*-neighbourhood $\mathcal{O}$ of μ, there is a w*-neighbourhood $\mathcal{O}' \subseteq \mathcal{O}$ such that the pair*

$$\big(\Omega(\mathcal{O}', \sigma_n), \Omega(\mathcal{O}, \sigma_n)\big)$$

is relatively δ-connected according to $d^{(V_n)}$ for all sufficiently large n.

These properties are both implied by the following.

3. *For every $\delta > 0$, μ has a base of w*-neighbourhoods $\mathcal{N}$ with the property that the set $\Omega(\mathcal{N}, \sigma_n)$ is δ-connected for all sufficiently large n.*

Proof. (1. $\Longrightarrow$ 2.) If property 2 does not hold, then there are some δ and $\mathcal{O}$ for which it fails. Let $\mathcal{O}_1 \supseteq \mathcal{O}_2 \supseteq \cdots$ be a base for the w*-topology at μ such that $\mathcal{O}_i \subseteq \mathcal{O}$ for every i. By the failure of property 2, there are integers $n_1 < n_2 < \ldots$ and pairs $\{\mathbf{x}_i, \mathbf{y}_i\} \subseteq \Omega(\mathcal{O}_i, \sigma_{n_i})$ for every i which cannot be connected by δ-paths that stay inside $\Omega(\mathcal{O}, \sigma_{n_i})$. This prevents the existence of δ-paths from $\mathbf{x}_i$ to $\mathbf{y}_i$ satisfying (3) in Definition 2.5.

(2. $\Longrightarrow$ 1.) Let $\mathbf{x}_i, \mathbf{y}_i \in \mathcal{X}^{V_{n_i}}$ be a pair of sequences as in Definition 2.5. Let $\mathcal{O}_1 \supseteq \mathcal{O}_2 \supseteq \ldots$ be a base for the w*-topology at μ.

For each j, property 2 gives a sub-neighbourhood $\mathcal{O}'_j \subseteq \mathcal{O}_j$ such that the pair

$$\big(\Omega(\mathcal{O}'_j, \sigma_n), \Omega(\mathcal{O}_j, \sigma_n)\big) \tag{4}$$

is relatively 2^{-j}-connected (according to $d^{(V_n)}$) for all sufficiently large n.

Since $P^{\sigma_{n_i}}_{\mathbf{x}_i}, P^{\sigma_{n_i}}_{\mathbf{y}_i} \overset{\text{weak}^*}{\longrightarrow} \mu$, there are $i_1 < i_2 < \ldots$ such that $\mathbf{x}_i, \mathbf{y}_i \in \Omega(\mathcal{O}'_j, \sigma_{n_i})$ for all $i \geqslant i_j$. Now the relative 2^{-j}-connectedness of the pair (4) implies that for each $i \geqslant i_j$ there is a 2^{-j}-path

$$\{\mathbf{x}_i = \mathbf{x}_{i,0},\ \mathbf{x}_{i,1},\ \ldots,\ \mathbf{x}_{i,\ell_i} = \mathbf{y}_i\} \subseteq \Omega(\mathcal{O}_j, \sigma_{n_i}).$$

These paths verify Definition 2.5 with

$$\delta_i := 2^{-\max\{j:\ i \geqslant i_j\}}.$$

(3. $\Longrightarrow$ 2.) For any $\mathcal{O}$, property 3 gives a sub-neighbourhood $\mathcal{N} \subseteq \mathcal{O}$ for which $\Omega(\mathcal{N}, \sigma_n)$ is δ-connected for all sufficiently large n, and this implies that the pair

$$\big(\Omega(\mathcal{N}, \sigma_n), \Omega(\mathcal{O}, \sigma_n)\big)$$

is relatively δ-connected. ∎

In all the examples of connected model spaces that I know, one actually has property 3 above, which is easier to verify. But I do not see a proof that these are equivalent, and I also do not know whether property 3 is isomorphism-invariant.

Remark. Definition 2.5 has a natural modification as follows. Let us say that $(\mathcal{X}^G, \mu, S, d)$ has **uniformly connected model spaces rel** Σ if there is a function $\ell : (0,1) \longrightarrow \mathbb{N}$ for which following holds:

> If $n_1 < n_2 < \ldots$, and $\mathbf{x}_i, \mathbf{y}_i \in \mathcal{X}^{V_{n_i}}$ are two sequences satisfying
>
> $$P^{\sigma_{n_i}}_{\mathbf{x}_i}, P^{\sigma_{n_i}}_{\mathbf{y}_i} \stackrel{\text{weak}^*}{\longrightarrow} \mu,$$
>
> then for every $\delta > 0$ and all sufficiently large i there are δ-paths
>
> $$\{\mathbf{x}_i = \mathbf{x}_{i,0},\ \mathbf{x}_{i,1},\ \ldots,\ \mathbf{x}_{i,\ell(\delta)} = \mathbf{y}_i\} \subseteq \mathcal{X}^{V_{n_i}}$$
>
> such that for any w*-neighbourhood $\mathcal{O}$ of μ we have
>
> $$\{\mathbf{x}_{i,j} :\ 0 \leqslant j \leqslant \ell(\delta)\} \subseteq \Omega(\mathcal{O}, \sigma_{n_i})$$
>
> for all sufficiently large i.

That is, one requires that the lengths of the δ-paths depend only on δ, not on n. It is easily shown that this variant is formally stronger than Definition 2.5

In the case of our principal examples, Bernoulli shifts, the proof below actually shows that this stronger property holds. I do not know of any examples that have connected but not uniformly connnected model spaces. $\lhd$

3 Model-surjective factor maps and Theorem A

As promised in the Introduction, for some groups G the Popa factor does not preserve the connectedness of model spaces. We now introduce a special kind of factor map which does always preserve this property, and show that all isomorphisms are factor maps of this kind. They are defined in terms of good models for the associated graphical joinings, but we will prove an equivalent characterization in terms of AL approximations to the factor map. We finish the subsection by proving Theorem A.

Later we show that complemented factor maps of Bernoulli shifts are also of this kind, which leads to the proof of Corollary D″.

Let

$$\Phi = \varphi^G : (\mathcal{X}^G, \mu, S, d_{\mathcal{X}}) \longrightarrow (\mathcal{Y}^G, \nu, S, d_{\mathcal{Y}})$$

be a factor map of metric G-processes, and let

$$\lambda = \int_{\mathcal{X}^G} \delta_{(x,\Phi(x))}\, \mu(\mathrm{d}x)$$

be the associated graphical joining of μ and ν. Also, let d be the Hamming average of the metrics $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$ on $\mathcal{X} \times \mathcal{Y}$.

Definition 3.1 (Model-surjective factor maps rel Σ)**.** *This factor map Φ is **model-surjective rel** Σ if, whenever $n_1 < n_2 < \ldots$ and $\mathbf{y}_i \in \mathcal{Y}^{V_{n_i}}$ is a sequence satisfying*

$$P_{\mathbf{y}_i}^{\sigma_{n_i}} \overset{\text{weak}^*}{\longrightarrow} \nu,$$

there is another sequence $\mathbf{x}_i \in \mathcal{X}^{V_{n_i}}$ such that

$$P_{(\mathbf{x}_i,\mathbf{y}_i)}^{\sigma_{n_i}} \overset{\text{weak}^*}{\longrightarrow} \lambda,$$

where this last convergence refers to the w-topology arising from the product topology on $\mathcal{X}^G \times \mathcal{Y}^G$.*

Clearly the sequence $\mathbf{x}_i$ obtained above must also satisfy $P_{\mathbf{x}_i}^{\sigma_n} \overset{\text{weak}^*}{\longrightarrow} \mu$. The idea behind Definition 3.1 is that, if the joint empirical distribution converges to λ, then this forces $\mathbf{x}_i$ to 'resemble' a pre-image of $\mathbf{y}_i$ under some model-space approximation to the factor map Φ.

The need to consider a sofic sub-approximation $(\sigma_{n_i})_{i\geqslant 1}$ of Σ is similar to the case of Definition 2.5.

Beware that at this point, Definition 3.1 requires a particular choice of the metrics $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$. Once we have shown that isomorphisms are model-surjective (Proposition 3.6 below), it will follow that Definition 3.1 actually depends only on the measure theoretic structure of Φ as a factor map from $(\mathcal{X}^G, \mu, S)$ to $(\mathcal{Y}^G, \nu, S)$, and moreover that it can be extended unambiguously to factor maps of general G-systems (Corollary 3.7).

First we need a simple lemma and corollary which relate the metrics $d^{(V_n)}$ and certain empirical distributions.

Lemma 3.2. *Let Σ be as before, and let $(\mathcal{X}, d)$ be a compact metric space. For any two sequences $\mathbf{x}_n, \mathbf{z}_n \in \mathcal{X}^{V_n}$, the following are equivalent:*

1. *we have* $d^{(V_n)}(\mathbf{x}_n, \mathbf{z}_n) \longrightarrow 0$*;*

2. *in the w* topology, any subsequential limit of the sequence*

$$P^{\sigma_n}_{(\mathbf{x}_n, \mathbf{z}_n)} \in \mathrm{Prob}(\mathcal{X}^G \times \mathcal{X}^G)$$

is supported on the diagonal

$$\{(x, x) : \ x \in \mathcal{X}^G\}.$$

In case $P^{\sigma_n}_{\mathbf{x}_n}$ *w*-converges to some* $\mu \in \mathrm{Prob}(\mathcal{X}^G)$*, either of the above conditions implies that* $P^{\sigma_n}_{\mathbf{z}_n}$ *w*-converges to the same limit.*

Proof. On $\mathcal{X}^G \times \mathcal{X}^G$, consider the function

$$F(x, x') := d(x_e, x'_e) \quad \text{for } x = (x_g)_g, \ x' = (x'_g)_g \in \mathcal{X}^G.$$

It is continuous, and a simple calculation gives

$$\int F \, \mathrm{d}P^{\sigma_n}_{(\mathbf{x}_n, \mathbf{z}_n)} = d^{(V_n)}(\mathbf{x}_n, \mathbf{y}_n).$$

Assuming condition (1), it follows that any subsequential limit

$$\lambda = \lim_{j \longrightarrow \infty} P^{\sigma_{n_j}}_{(\mathbf{x}_{n_j}, \mathbf{z}_{n_j})}$$

must satisfy $\int F \, \mathrm{d}\lambda = 0$, hence be supported on the set $\{(x, x') : \ x_e = x'_e\}$. Since λ is S-invariant (see [1, Lemma 3.2]), it must be supported on the diagonal.

On the other hand, if (1) fails, then there are some $\delta > 0$ and some subsequence $n_1 < n_2 < \dots$ for which

$$d^{(V_{n_j})}(\mathbf{x}_{n_j}, \mathbf{z}_{n_j}) \longrightarrow \delta \quad \text{as } j \longrightarrow \infty.$$

By the sequential compactness of the w*-topology on $\mathrm{Prob}(\mathcal{X}^G \times \mathcal{X}^G)$, there is a further subsequence for which the empirical distributions converge to some λ. This λ must then satisfy $\int F \, \mathrm{d}\lambda = \delta > 0$, and so cannot be supported on the diagonal.

Finally, any measure supported on the diagonal must have equal first and second marginals, so condition (2) clearly implies the last part of the lemma. □

Corollary 3.3. *Let*

$$\Phi = \varphi^G : (\mathcal{X}^G, \mu, S, d_{\mathcal{X}}) \longrightarrow (\mathcal{Y}^G, \nu, S, d_{\mathcal{Y}})$$

be a factor map of metric G-processes, and let

$$\lambda = \int_{\mathcal{X}^G} \delta_{(x,\Phi(x))}\, \mu(\mathrm{d}x)$$

be the resulting graphical joining of these two systems. Suppose that $\mathbf{x}_n \in \mathcal{X}^{V_n}$ and $\mathbf{y}_n$, $\mathbf{w}_n \in \mathcal{Y}^{V_n}$ are sequence such that

$$P^{\sigma_n}_{(\mathbf{x}_n,\mathbf{y}_n)}, P^{\sigma_n}_{(\mathbf{x}_n,\mathbf{w}_n)} \stackrel{\text{weak}^*}{\longrightarrow} \lambda.$$

Then $d_{\mathcal{Y}}^{(V_n)}(\mathbf{y}_n, \mathbf{w}_n) \longrightarrow 0$.

Proof. For each n, let

$$\theta_n := P^{\sigma_n}_{(\mathbf{x}_n,\mathbf{y}_n,\mathbf{w}_n)} \in \mathrm{Prob}((\mathcal{X} \times \mathcal{Y} \times \mathcal{Y})^G),$$

and let

$$\theta := \lim_{j \longrightarrow \infty} \theta_{n_j}$$

be any w*-subsequential limit of these measures.

The projections of θ_n onto the two copies of $(\mathcal{X} \times \mathcal{Y})^G$ are $P^{\sigma_n}_{(\mathbf{x}_n,\mathbf{y}_n)}$ and $P^{\sigma_n}_{(\mathbf{x}_n,\mathbf{w}_n)}$, so we know that the limiting measure θ has both of these projections equal to λ. Hence θ is supported on

$$\{(x, y, y') :\ x \in \mathcal{X}^G,\ y = \varphi^G(x) = y'\}.$$

Therefore the limit of $P^{\sigma_n}_{(\mathbf{y}_{n_j},\mathbf{w}_{n_j})}$ is supported on the diagonal in $(\mathcal{Y} \times \mathcal{Y})^G$, and Lemma 3.2 completes the proof. □

Now we can begin the study of model-surjectivity. Graphical joinings give the easiest way to define this property, but it is useful to have a more 'functional' characterization. This can be given in terms of AL approximations to the factor map.

Lemma 3.4. *For $\Phi = \varphi^G$ as above, the following are equivalent.*

1. *The factor map Φ is model-surjective rel Σ.*

2. *Suppose that* $\psi_k \xrightarrow{\text{aL}} \varphi$ *rel* $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$. *If* $n_1 < n_2 < \dots$ *and* $\mathbf{y}_i \in \mathcal{Y}^{V_{n_i}}$ *is a sequence satisfying*

$$P^{\sigma_{n_i}}_{\mathbf{y}_i} \xrightarrow{\text{weak}^*} \nu,$$

then there is a sequence $\mathbf{x}_i \in \mathcal{X}^{V_{n_i}}$ *such that we have*

$$P^{\sigma_{n_i}}_{\mathbf{x}_i} \xrightarrow{\text{weak}^*} \mu$$

and

$$d^{(V_{n_i})}_{\mathcal{Y}}\big(\mathbf{y}_i, \psi^{\sigma_{n_i}}_{k_i}(\mathbf{x}_i)\big) \longrightarrow 0$$

whenever $k_1 \leqslant k_2 \leqslant \dots$ *grows sufficiently slowly.*

Proof. Let $\psi_k \xrightarrow{\text{aL}} \varphi$ be an AL approximating sequence rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$. Also, let $\xi : \mathcal{X}^G \longrightarrow \mathcal{X}$ be the projection onto the $\{e\}$-coordinate, so $\xi^G = \mathrm{id}_{\mathcal{X}^G}$. An easy check (or see [1, Corollary 4.7]) gives that

$$(\xi, \psi_k) \xrightarrow{\text{aL}} (\xi, \varphi) \quad \text{rel } (\mu, d_{\mathcal{X}}, d), \tag{5}$$

where d is the Hamming average of $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$ on $\mathcal{X} \times \mathcal{Y}$, and (ξ, φ) denotes the map

$$\mathcal{X}^G \longrightarrow \mathcal{X} \times \mathcal{Y} : x \mapsto (\xi(x), \varphi(x)).$$

Let $n_1 < n_2 < \dots$, and let $\mathbf{y}_i \in \mathcal{Y}^{V_{n_i}}$ be a sequence whose empirical distributions tend to ν. By re-labeling the sofic sub-approximation $(\sigma_{n_i})_{i \geqslant 1}$ if necessary, we may assume that $n_i = i$ for all i, and hence write $\mathbf{y}_n$ as an element of $\mathcal{Y}^{V_n}$.

Now suppose $\mathbf{x}_n \in \mathcal{X}^{V_n}$ is another sequence whose empirical distributions tend to μ. By (5) and Corollary 2.4, it follows that

$$P^{\sigma_n}_{(\mathbf{x}_n, \psi^{\sigma_n}_{k_n}(\mathbf{x}_n))} = P^{\sigma_n}_{(\xi, \psi_{k_n})^{\sigma_n}(\mathbf{x}_n)} \xrightarrow{\text{weak}^*} \lambda,$$

provided the sequence $k_1 \leqslant k_2 \leqslant \dots$ grows sufficiently slowly.

The result follows because Lemma 3.2 (respectively Corollary 3.3) shows that

$$P^{\sigma_n}_{(\mathbf{x}_n, \mathbf{y}_n)} \xrightarrow{\text{weak}^*} \lambda$$

if (respectively only if)

$$d^{(V_n)}\big((\mathbf{x}_n, \mathbf{y}_n), (\mathbf{x}_n, \psi^{\sigma_n}_{k_n}(\mathbf{x}_n))\big) = \frac{1}{2} d^{(V_n)}_{\mathcal{Y}}\big(\mathbf{y}_n, \psi^{\sigma_n}_{k_n}(\mathbf{x}_n)\big) \longrightarrow 0.$$

□

Proposition 3.5. *If*

$$\Phi : (\mathcal{X}^G, \mu, S, d_{\mathcal{X}}) \longrightarrow (\mathcal{Y}^G, \nu, S, d_{\mathcal{Y}})$$
$$\text{and} \quad \widetilde{\Phi} : (\mathcal{Y}^G, \nu, S, d_{\mathcal{Y}}) \longrightarrow (\mathcal{Z}^G, \theta, S, d_{\mathcal{Z}})$$

are both model-surjective rel Σ, then so is their composition.

Proof. As in the proof above, after passing to a sofic sub-approximation we may assume that $\mathbf{z}_n \in \mathcal{Z}^{V_n}$ is a sequence whose empirical distributions tend to θ. By the two assumed instances of model-surjectivity rel Σ, we may find first a sequence $\mathbf{y}_n \in \mathcal{Y}^{V_n}$ and then a sequence $\mathbf{x}_n \in \mathcal{X}^{V_n}$ such that

$$P^{\sigma_n}_{(\mathbf{x}_n, \mathbf{y}_n)} \overset{\text{weak}^*}{\longrightarrow} \lambda \quad \text{and} \quad P^{\sigma_n}_{(\mathbf{y}_n, \mathbf{z}_n)} \overset{\text{weak}^*}{\longrightarrow} \widetilde{\lambda},$$

where λ and $\widetilde{\lambda}$ are the graphical joinings associated to Φ and $\widetilde{\Phi}$ respectively.

Now let

$$\lambda_1 = \lim_{j \longrightarrow \infty} P^{\sigma_{n_j}}_{(\mathbf{x}_{n_j}, \mathbf{y}_{n_j}, \mathbf{z}_{n_j})}$$

be any subsequential limit of the triple empirical distributions. Since its projection onto the first two coordinates must be λ and its projection onto the second two coordinates must be $\widetilde{\lambda}$, it is supported on the set

$$\big\{\big(x, \Phi(x), \widetilde{\Phi}(\Phi(x))\big) : \ x \in \mathcal{X}^G\big\},$$

and hence it must be the full graphical joining associated to $(\Phi, \widetilde{\Phi} \circ \Phi)$. Therefore any subsequence of

$$P^{\sigma_n}_{(\mathbf{x}_n, \mathbf{z}_n)}$$

converges to the graphical joining associated to $\widetilde{\Phi} \circ \Phi$, as required. □

Proposition 3.6. *If $\Phi = \varphi^G$ is an isomorphism, then it is model-surjective for any sofic approximation Σ.*

Proof. Let $\widetilde{\varphi}$ be such that $\Phi^{-1} = \widetilde{\varphi}^G$, let $\xi : \mathcal{Y}^G \longrightarrow \mathcal{Y}$ be the projection to the $\{e\}$-indexed coordinate, and let $\widetilde{\psi}_m \overset{\text{aL}}{\longrightarrow} \widetilde{\varphi}$ rel $(\nu, d_{\mathcal{Y}}, d_{\mathcal{X}})$. Let λ be the graphical joining of Φ. Then we also have

$$(\widetilde{\psi}_m, \xi) \overset{\text{aL}}{\longrightarrow} (\widetilde{\varphi}, \xi) \quad \text{rel } (\nu, d_{\mathcal{Y}}, d) \tag{6}$$

(see again [1, Corollary 4.7] for a careful proof of this).

As in the proofs above, to show the model-surjectivity of Φ we may pass to a sofic sub-approximation, and so suppose that $\mathbf{y}_n \in \mathrm{Prob}(\mathcal{Y}^{V_n})$ is a sequence whose empirical distributions converge to ν. Then (6) and Lemma 2.1 give that

$$P^{\sigma_n}_{\left(\tilde{\psi}^{\sigma_n}_{m_n}(\mathbf{y}_n),\mathbf{y}_n\right)} \overset{\text{weak}^*}{\longrightarrow} \lambda$$

provided $(m_n)_{n\geqslant 1}$ grows sufficiently slowly.

Letting $\mathbf{x}_n := \tilde{\psi}^{\sigma_n}_{m_n}(\mathbf{y}_n)$, this completes the proof. □

The first important consequence of Propositions 3.5 and 3.6 is the following.

Corollary 3.7. *If*

$$\begin{array}{ccc} (\mathcal{X}^G, \mu, S, d_{\mathcal{X}}) & \xrightarrow{\cong} & (\mathcal{X}_1^G, \mu_1, S, d_{\mathcal{X}_1}) \\ \Big\downarrow{\scriptstyle \Phi} & & \Big\downarrow{\scriptstyle \Psi} \\ (\mathcal{Y}^G, \nu, S, d_{\mathcal{Y}}) & \xrightarrow{\cong} & (\mathcal{Y}_1^G, \nu_1, S, d_{\mathcal{Y}_1}) \end{array}$$

is a commutative square of factor maps in which the horizontal arrows are isomorphisms and Φ is model-surjective rel Σ, then Ψ is also model-surjective rel Σ.

In particular, model-surjectivity rel Σ in Definition 3.1 is independent of the choice of generating metrics $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$. □

We now have all the ingredients needed to prove Theorem A.

Proof of Theorem A. Let $n_1 < n_2 < \ldots$, and let $\mathbf{y}_i, \mathbf{w}_i \in \mathcal{Y}^{V_{n_i}}$ be sequences such that

$$P^{\sigma_{n_i}}_{\mathbf{y}_i}, P^{\sigma_{n_i}}_{\mathbf{w}_i} \overset{\text{weak}^*}{\longrightarrow} \nu.$$

We must show that these pairs may be connected by $o(1)$-paths consisting of good models, as in Definition 2.5. By passing to the sofic sub-approximation $(\sigma_{n_i})_{i\geqslant 1}$, we may relabel all these sequences and so assume that $i = n_i$ for all i, and hence write the index as n itself.

Now let $\mathbf{x}_n, \mathbf{z}_n \in \mathcal{X}^{V_n}$ be the sequences given by the model-surjectivity of Φ applied to the sequences $\mathbf{y}_n$ and $\mathbf{w}_n$, respectively. Since $(\mathcal{X}^G, \mu, S, d_{\mathcal{X}})$ has connected model space rel Σ, there are parameters $\delta_n \downarrow 0$ and a sequence of δ_n-paths

$$\mathbf{x}_n = \mathbf{x}_{n,0}, \mathbf{x}_{n,1}, \ldots, \mathbf{x}_{n,\ell_n} = \mathbf{z}_n$$

which are eventually contained in $\Omega(\mathcal{O}, \sigma_n)$ for any w*-neighbourhood $\mathcal{O}$ of μ.

Also, let $\psi_k \xrightarrow{\mathrm{aL}} \varphi$ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$. Corollary 2.3 gives parameters $\varepsilon_k \downarrow 0$ and $K_k < \infty$, a sequence of w*-neighbourhoods $\mathcal{O}_1 \supseteq \mathcal{O}_2 \supseteq \dots$ of μ, and a base of w*-neighbourhoods $\mathcal{N}_1 \supseteq \mathcal{N}_2 \supseteq \dots$ at ν, such that for every k we have both

$$\psi_k^{\sigma_n} | \Omega(\mathcal{O}_k, \sigma_n) \quad \text{is } \varepsilon_k\text{-almost } K_k\text{-Lipschitz for all sufficiently large } n \tag{7}$$

and

$$\psi_k^{\sigma_n}\big(\Omega(\mathcal{O}_k, \sigma_n)\big) \subseteq \Omega(\mathcal{N}_k, \sigma_n) \quad \text{for all sufficiently large } n. \tag{8}$$

Now choose a sequence $k_1 \leqslant k_2 \leqslant \dots$ growing so slowly that all of the following hold:

i) we have $\varepsilon_{k_n} + K_{k_n}\delta_n \longrightarrow 0$ as $n \longrightarrow \infty$ (this is possible because $\varepsilon_k \longrightarrow 0$ and $\delta_n \longrightarrow 0$);

ii) we have
$$\{\mathbf{x}_{n,0}, \mathbf{x}_{n,1}, \dots, \mathbf{x}_{n,\ell_n}\} \subseteq \Omega(\mathcal{O}_{k_n}, \sigma_n)$$
for all sufficiently large n;

iii) we have
$$d_{\mathcal{Y}}^{(V_n)}\big(\mathbf{y}_n, \psi_{k_n}^{\sigma_n}(\mathbf{x}_n)\big),\ d_{\mathcal{Y}}^{(V_n)}\big(\mathbf{w}_n, \psi_{k_n}^{\sigma_n}(\mathbf{z}_n)\big) \longrightarrow 0 \quad \text{as } n \longrightarrow \infty,$$
as promised by condition 2 in Lemma 3.4.

iv) we have
$$\psi_{k_n}^{\sigma_n} | \Omega(\mathcal{O}_{k_n}, \sigma_n) \quad \text{is } \varepsilon_{k_n}\text{-almost } K_{k_n}\text{-Lipschitz}$$
for all sufficiently large n, as is possible by (7)

v) we have
$$\psi_{k_n}^{\sigma_n}\big(\Omega(\mathcal{O}_{k_n}, \sigma_n)\big) \subseteq \Omega(\mathcal{N}_{k_n}, \sigma_n)$$
for all sufficiently large n, as is possible by (8).

Finally, define
$$\mathbf{y}_{n,i} := \psi_{k_n}^{\sigma_n}(\mathbf{x}_{n,i})$$
for each n and each $i = 0, 1, \dots, \ell_n$.

By properties (ii) and (iv), the sequence

$$\mathbf{y}_{n,0}, \mathbf{y}_{n,1}, \dots, \mathbf{y}_{n,\ell_n}$$

is a $(\varepsilon_{k_n} + K_{k_n}\delta_n)$-path according to $d_{\mathcal{Y}}^{(V_n)}$ for all sufficiently large n. Combined with properties (i) and (iii), this shows that

$$\mathbf{y}_n, \mathbf{y}_{n,0}, \mathbf{y}_{n,1}, \ldots, \mathbf{y}_{n,\ell_n}, \mathbf{w}_n$$

is a δ'_n-path from $\mathbf{y}_n$ to $\mathbf{w}_n$ for some sequence of parameters $\delta'_n \downarrow 0$.

Finally, properties (ii) and (v) imply that these image paths are eventually contained in $\Omega(\mathcal{N}, \sigma_n)$ for every w*-neighbourhood $\mathcal{N}$ of ν, because $\mathcal{N}_1 \supseteq \mathcal{N}_2 \supseteq \ldots$ is a base at ν. □

4 Connected model spaces for inverse limits

This section gives the proof of Theorem B.

Proof of Theorem B. Up to isomorphism, the setting of this theorem may be represented as follows. Let $(\mathcal{X}_i^G, \mu_i, S)$, $i = 1, 2, \ldots$ be an infinite sequence of G-processes, and let

$$\lambda \in \mathrm{Prob}(\mathcal{X}_1^G \times \mathcal{X}_2^G \times \cdots)$$

be a joining of all of them. Let $\pi_k : \prod_{i\geqslant 1} \mathcal{X}_i \longrightarrow \prod_{i=1}^k \mathcal{X}_i$ be the coordinate projection for each k, and let $\lambda_k := \pi_{k*}^G \lambda$. We assume that the G-system

$$\Big(\prod_{i=1}^k \mathcal{X}_i^G, \lambda_k, S\Big)$$

has connected model spaces rel Σ for each k, and must prove the same for the infinite joining λ. This is easiest using the reformulation in condition 2 of Proposition 2.6.

For each i, let $d_{\mathcal{X}_i}$ be a compact generating metric of diameter at most 1 for the space $\mathcal{X}_i$. Define metrics d_k on $\prod_{i=1}^k \mathcal{X}_i$ and d on $\prod_{i\geqslant 1} \mathcal{X}_i$ by

$$d_k\big((x_i)_{i=1}^k, (x'_i)_{i=1}^k\big) := \sum_{i=1}^k 2^{-i} d_{\mathcal{X}_i}(x_i, x'_i)$$

and

$$d\big((x_i)_{i\geqslant 1}^k, (x'_i)_{i\geqslant 1}\big) := \sum_{i\geqslant 1} 2^{-i} d_{\mathcal{X}_i}(x_i, x'_i),$$

so these generate the compact product topologies on their respective spaces.

Now suppose that $\mathcal{O}$ is a w*-neighbourhood of λ. By shrinking it if necessary, we may assume that it has the form

$$\big\{\theta :\ \pi^G_{k*}\theta \in \mathcal{O}_1\big\}$$

for some $k \in \mathbb{N}$ and some w*-neighbourhood $\mathcal{O}_1$ of λ_k, since sets of this form are a base of w*-neighbourhoods around λ. Under this assumption, we obtain also

$$\begin{aligned}\Omega(\mathcal{O}, \sigma_n) &= \Big\{\mathbf{x} \in \Big(\prod_{i\geqslant 1} \mathcal{X}_i\Big)^{V_n} :\ \pi^{\sigma_n}_k(\mathbf{x}) \in \Omega(\mathcal{O}_1, \sigma_n)\Big\} \\ &= \Omega(\mathcal{O}_1, \sigma_n) \times \Big(\prod_{i\geqslant n+1} \mathcal{X}_i\Big)^{V_n} \quad \text{for each } n. \quad (9)\end{aligned}$$

Since the process defined by λ_k has connected model spaces rel Σ, Proposition 2.6 gives another w*-neighbourhood $\mathcal{O}'_1 \subseteq \mathcal{O}_1$ of λ_k such that the pair

$$\big(\Omega(\mathcal{O}'_1, \sigma_n), \Omega(\mathcal{O}_1, \sigma_n)\big)$$

is relatively δ-connected according to $d_k^{(V_n)}$ for all sufficiently large n.

Let

$$\mathcal{O}' := \big\{\theta :\ \pi^G_{k*}\theta \in \mathcal{O}'_1\big\} \subseteq \mathcal{O}.$$

Equation (9) has an obvious analog for $\mathcal{O}'$ and $\mathcal{O}'_1$. Combining these, it follows that the pair

$$\big(\Omega(\mathcal{O}', \sigma_n), \Omega(\mathcal{O}, \sigma_n)\big)$$

is also relatively δ-connected according to $d^{(V_n)}$ for all sufficiently large n. This verifies condition 2 in Proposition 2.6 for the process defined by λ. □

5 Connected model spaces for Bernoulli systems

This subsection proves Theorem C. We actually prove the slightly stronger property (3) from Proposition 2.6. Let $(\mathcal{X}, d)$ be a compact metric space of diameter at most 1, and let $\nu \in \mathrm{Prob}(\mathcal{X})$.

This section makes several simple appeals to the phenomenon of measure concentration for product measures and Hamming metrics: see, for instance, [16] for a dedicated exposition and [8, Chapter $3\frac{1}{2}$] for a geometrically-flavoured overview. The specific result that we need is the following: see [16, Corollary 1.17].

Proposition 5.1. *For any $\delta > 0$ there is a $\beta > 0$ such that*

$$\sup_f \nu^{\times n}\Big\{\Big|f - \int f \,\mathrm{d}\nu^{\times n}\Big| \geqslant \delta\Big\} = O(\mathrm{e}^{-\beta n})$$

for all $n \geqslant 1$, where the supremum is over all functions $f : \mathcal{X}^n \longrightarrow \mathbb{R}$ which are 1-Lipschitz for the normalized Hamming distance $d^{(n)}$. □

Next we identify certain special neighbourhoods of the product measure $\nu^{\times G}$. Let F be a finite subset of G, and for each $D \subseteq F$ define an operator E_D on $C(\mathcal{X}^F)$ as follows. First, if $F = \{g_1, \ldots, g_m\}$ and $D = F\backslash\{g_j\}$ for some $j \leqslant m$, then

$$E_D f(x) := \int_X f(x_{g_1}, \ldots, x_{g_{j-1}}, y, x_{g_{j+1}}, \ldots, x_{g_m})\, \nu(\mathrm{d}y).$$

In general, if $D = F\backslash\{g_{j_1}, \ldots, g_{j_k}\}$, then

$$E_D := E_{F\backslash\{g_{j_1}\}} \circ \cdots \circ E_{F\backslash\{g_{j_k}\}},$$

where the order of this composition is unimportant. This E_D is the operator of conditional expectation with respect to $\nu^{\times F}$ onto the σ-algebra of Borel subsets of $\mathcal{X}^F$ that depend only on coordinates in D. It follows that

$$\int f \,\mathrm{d}\nu^{\times F} = \int E_D f \,\mathrm{d}\nu^{\times F} \quad \forall f \in C(\mathcal{X}^F). \tag{10}$$

Now suppose that $\mathcal{F} \subseteq C(\mathcal{X}^G)$ is a family of F-local functions. Let us call it **hereditary** if every member of $\mathcal{F}$ is 1-Lipschitz according to $d_{\mathcal{X}}^{(F)}$ and if

$$f \in \mathcal{F} \quad \Longrightarrow \quad E_D f \in \mathcal{F} \quad \forall D \subseteq F.$$

It is easily checked that each E_D preserves the property of being 1-Lipschitz with respect to $d_{\mathcal{X}}^{(F)}$. A **hereditary neighbourhood** of $\nu^{\times G}$ is a w*-neighbourhood of the form

$$\mathcal{O} := \Big\{\theta \in \mathrm{Prob}(\mathcal{X}^G) : \int f \,\mathrm{d}\theta \approx_\varepsilon \int f \,\mathrm{d}\nu^{\times G} \;\forall f \in \mathcal{F}\Big\}. \tag{11}$$

for some finite $F \subseteq G$, some $\varepsilon > 0$, and some finite hereditary family $\mathcal{F}$ of F-local continuous functions.

Proposition 5.2. *If $\mathcal{O}$ is a hereditary neighbourhood of $\nu^{\times G}$ and $\delta > 0$, then the set $\Omega(\mathcal{O}, \sigma_n)$ is δ-connected according to $d^{(V_n)}$ for all sufficiently large n.*

We will see the importance of assuming that $\mathcal{O}$ is hereditary during the course of the proof.

Proof of Theorem C from Proposition 5.2. Hereditary neighbourhoods form a base for the w*-topology at μ. Therefore Proposition 5.2 verifies condition 3 in Proposition 2.6, which implies connected model spaces. □

Proposition 5.2 will be proved by showing how any pair of points in $\Omega(\mathcal{O}, \sigma_n)$ may be connected by a 'random' δ-path, provided n is sufficiently large.

We insert randomness into the proof as follows. Fix $\kappa \in (0, 1)$, to be specified later. For each n, let $(\boldsymbol{\xi}_{n,t})_{t\geqslant 0}$ be a discrete-time random walk on $\mathcal{X}^{V_n}$ with the following transition probabilities:

$$\mathbb{P}(\boldsymbol{\xi}_{n,t+1} \in \cdot \,|\, \boldsymbol{\xi}_{n,t}) = \bigtimes_{v\in V_n} \big(\kappa\delta_{\xi_{n,t,v}} + (1-\kappa)\nu\big), \tag{12}$$

where we write $\boldsymbol{\xi}_{n,t} = (\xi_{n,t,v})_{v\in V_n}$. Thus, $\boldsymbol{\xi}_{n,t+1}$ is obtained from $\boldsymbol{\xi}_{n,t}$ by considering each coordinate in $\mathcal{X}^{V_n}$ independently, and either re-sampling it from the distribution ν with probability $1-\kappa$, or leaving it unchanged with probability κ. Let $\mathbb{P}_n^{\mathbf{x}}$ be the law of $(\boldsymbol{\xi}_{n,t})_{t\geqslant 0}$ on $\mathcal{X}^{V_n} \times \mathcal{X}^{V_n} \times \cdots$ conditioned on starting from $\boldsymbol{\xi}_{n,0} = \mathbf{x}$, and let $\mathbb{E}_n^{\mathbf{x}}$ denote expectation with respect to $\mathbb{P}_n^{\mathbf{x}}$.

We will find a δ-path between two good models $\mathbf{x}$ and $\mathbf{y}$ by starting a copy of this random walk at each of $\mathbf{x}$ and $\mathbf{y}$, and showing that after a certain bounded time the following hold:

(i) these random walks have probably stayed inside the set of good models,

(ii) they have probably taken only steps smaller than δ in the Hamming distance, and

(iii) they can be coupled in such a way that with high probability they end up close to each other.

We break the necessary estimates into three separate lemmas.

Lemma 5.3. *For any hereditary neighbourhood $\mathcal{O}$ and any $t \in \mathbb{N}$, we have*

$$\inf_{\mathbf{x}\in\Omega(\mathcal{O},\sigma_n)} \mathbb{P}_n^{\mathbf{x}}\big\{\boldsymbol{\xi}_{n,t} \in \Omega(\mathcal{O}, \sigma_n)\big\} \longrightarrow 1 \quad \textit{as } n \longrightarrow \infty.$$

Proof. Let $\mathcal{O}$ be as in (11) for some $\varepsilon > 0$ and some hereditary family $\mathcal{F}$ of F-local functions.

Since $\mathcal{F}$ is finite, it suffices to show that for any one $f \in \mathcal{F}$ we have

$$\inf_{\mathbf{x}\in\Omega(\mathcal{O},\sigma_n)} \mathbb{P}_n^{\mathbf{x}}\Big\{ \int f \,\mathrm{d}P^{\sigma_n}_{\boldsymbol{\xi}_{n,t}} \approx_\varepsilon \int f \,\mathrm{d}\nu^{\times G}\Big\} \longrightarrow 1.$$

We break this estimate into two further steps.

Step 1. Observe that

$$\mathbb{E}_n^{\mathbf{x}} \int f \,\mathrm{d}P^{\sigma_n}_{\boldsymbol{\xi}_{n,t}} = \frac{1}{|V_n|}\sum_{v\in V_n} \mathbb{E}_n^{\mathbf{x}} f\big(\Pi_v^{\sigma_n}(\boldsymbol{\xi}_{n,t})\big) = \frac{1}{|V_n|}\sum_{v\in V_n} \mathbb{E}_n^{\mathbf{x}} f\big((\boldsymbol{\xi}_{n,t,\sigma_n^g(v)})_{g\in F}\big).$$

Since F is finite and Σ is a sofic approximation, as $n \longrightarrow \infty$ it holds with high probability in the choice of $v \in V_n$ that

$$\text{the points } \sigma_n^g(v) \text{ for } g \in F \text{ are distinct.} \tag{13}$$

For such v, iterating the formula (12) gives

$$\mathbb{E}_n^{\mathbf{x}} f\big((\boldsymbol{\xi}_{n,t,\sigma_n^g(v)})_{g\in F}\big) = \sum_{D\subseteq F} \kappa^{t|D|}(1-\kappa^t)^{|F|-|D|} E_D f(\Pi_v^{\sigma_n}(\mathbf{x})).$$

Let D be a random subset of F which contains each element of F independently with probability κ^t. Let $\mathbb{P}'$ be its law and let $\mathbb{E}'$ be expectation with respect to $\mathbb{P}'$. Then the above leads to

$$\mathbb{E}_n^{\mathbf{x}} \int f \,\mathrm{d}P^{\sigma_n}_{\boldsymbol{\xi}_{n,t}} = \mathbb{E}'\Big(\frac{1}{|V_n|}\sum_{v\in V} E_D f(\Pi_v^{\sigma_n}(\mathbf{x}))\Big) + o(1),$$

where the $o(1)$-correction results from those few vertices v which fail the requirement (13).

Therefore

$$\begin{aligned}
&\Big|\mathbb{E}_n^{\mathbf{x}} \int f \,\mathrm{d}P^{\sigma_n}_{\boldsymbol{\xi}_{n,t}} - \int f \,\mathrm{d}\nu^{\times G}\Big| \\
&\leqslant \mathbb{E}'\Big|\frac{1}{|V_n|}\sum_{v\in V} E_D f(\Pi_v^{\sigma_n}(\mathbf{x})) - \int f \,\mathrm{d}\nu^{\times G}\Big| + o(1) \\
&= \mathbb{E}'\Big|\frac{1}{|V_n|}\sum_{v\in V} E_D f(\Pi_v^{\sigma_n}(\mathbf{x})) - \int E_D f \,\mathrm{d}\nu^{\times G}\Big| + o(1) \\
&= \mathbb{E}'\Big|\int E_D f \,\mathrm{d}(P^{\sigma_n}_{\mathbf{x}} - \nu^{\times G})\Big| + o(1) \\
&= \mathbb{E}'\Big(\Big|\int E_D f \,\mathrm{d}(P^{\sigma_n}_{\mathbf{x}} - \nu^{\times G})\Big| \cdot 1_{\{D\neq\varnothing\}}\Big) + o(1),
\end{aligned}$$

where the equality of second and third lines uses (10), and the last equality uses that $E_\varnothing f$ is constant and equal to $\int f \, \mathrm{d}\nu^{\times F}$. Since $\mathbf{x} \in \Omega(\mathcal{O}, \sigma_n)$ with $\mathcal{O}$ as in (11), the last line above is strictly bounded by

$$\mathbb{P}'\{D \neq \varnothing\} \cdot \varepsilon + o(1) = (1 - (1-\kappa^t)^{|F|})\varepsilon + o(1),$$

where the $o(1)$-correction depends only on the sofic approximation Σ and on $\|f\|_\infty$. Therefore this bound is strictly less than, say, $(1 - \frac{1}{2}(1-\kappa^t)^{|F|})\varepsilon$ for all sufficiently large n.

Step 2. On the other hand, we may regard the quantity $\int f \, \mathrm{d}P^{\sigma_n}_{\boldsymbol{\xi}_{n,t}}$ as a random variable on the probability space

$$\big(\mathcal{X}^{V_n}, \mathbb{P}^{\mathbf{x}}_n\{\boldsymbol{\xi}_{n,t} \in \cdot\}\big).$$

Using the initial condition $\boldsymbol{\xi}_{n,0} = \mathbf{x}$ and the transition probabilities (12), the probability measure here is equal to

$$\bigtimes_{v \in V_n} \big(\kappa^t \delta_{x_v} + (1-\kappa^t)\nu\big),$$

which is a product measure on $\mathcal{X}^{V_n}$. By the definition of the empirical distribution, and recalling that f is F-local and 1-Lipschitz according to $d^{(F)}_{\mathcal{X}}$, this random variable is an $|F|$-Lipschitz function on this product space. Therefore Proposition 5.1 gives some $\beta > 0$, depending only on the ratio $(1-\kappa^t)^{|F|}\varepsilon/|F|$, such that

$$\mathbb{P}^{\mathbf{x}}_n\Big\{\Big|\int f \, \mathrm{d}P^{\sigma_n}_{\boldsymbol{\xi}_{n,t}} - \mathbb{E}^{\mathbf{x}}_n \int f \, \mathrm{d}P^{\sigma_n}_{\boldsymbol{\xi}_{n,t}}\Big| \geqslant \frac{1}{2}(1-\kappa^t)^{|F|})\varepsilon\Big\} = O\big(\mathrm{e}^{-\beta|V_n|}\big).$$

Thus this probability tends to 0 as $n \longrightarrow \infty$ uniformly in $\mathbf{x}$.

Combining the estimates from Steps 1 and 2 completes the proof. □

Lemma 5.4. *Suppose that $1 - \kappa < \delta$. For any $t \in \mathbb{N} \cup \{0\}$, the random walks constructed above satisfy*

$$\inf_{\mathbf{x} \in \mathcal{X}^{V_n}} \mathbb{P}^{\mathbf{x}}_n\big\{d^{(V_n)}(\boldsymbol{\xi}_{n,t}, \boldsymbol{\xi}_{n,t+1}) < \delta\big\} \longrightarrow 1 \quad \textit{as } n \longrightarrow \infty.$$

Proof. By time-homogeneity, it suffices to prove this when $t = 0$.

Since $\mathrm{diam}(\mathcal{X}, d) \leqslant 1$, we always have

$$d^{(V_n)}(\mathbf{x}, \boldsymbol{\xi}_{n,1}) \leqslant \frac{1}{|V_n|} \sum_{v \in V_n} 1_{\{x_v \neq \xi_{n,1,v}\}}.$$

This is an average of indicator functions of independent events, all of them having probability at most $1-\kappa < \delta$. Therefore another appeal to Proposition 5.1 (or just the special case of a Chernoff bound) gives a $\beta > 0$ for which

$$\mathbb{P}_n^{\mathbf{x}}\big\{d^{(V_n)}(\mathbf{x}, \boldsymbol{\xi}_{n,1}) \geqslant \delta\big\} = O(\mathrm{e}^{-\beta|V_n|}).$$

Since the right-hand bound is independent of $\mathbf{x}$, this completes the proof. □

Lemma 5.5. *If $s \in \mathbb{N}$ is so large that $\kappa^s < \delta/4$, then the following holds. For any $n \in \mathbb{N}$ and any $\mathbf{x}, \mathbf{y} \in \mathcal{X}^{V_n}$, the distributions*

$$\mathbb{P}_n^{\mathbf{x}}\big\{(\boldsymbol{\xi}_{n,0}, \dots, \boldsymbol{\xi}_{n,s}) \in \cdot\big\} \quad \textit{and} \quad \mathbb{P}_n^{\mathbf{y}}\big\{(\boldsymbol{\zeta}_{n,0}, \dots, \boldsymbol{\zeta}_{n,s}) \in \cdot\big\}$$

of the random walks started at $\mathbf{x}$ and $\mathbf{y}$ up to the finite time-horizon s have a coupling $\mathbb{Q}$ such that

$$\mathbb{Q}\big\{d^{(V_n)}(\boldsymbol{\xi}_{n,s}, \boldsymbol{\zeta}_{n,s}) < \delta\big\} > 1/2. \tag{14}$$

Proof. By an induction on s using (12), we have

$$\mathbb{P}_n^{\mathbf{x}}\{\boldsymbol{\xi}_{n,s} \in \cdot\} = \bigtimes_{v \in V_n} \big(\kappa^s \delta_{x_v} + (1-\kappa^s)\nu\big),$$

and similarly for $\mathbb{P}_n^{\mathbf{y}}\{\boldsymbol{\zeta}_{n,s} \in \cdot\}$.

Consider a random triple $(\boldsymbol{\alpha}, \boldsymbol{\xi}, \boldsymbol{\zeta})$ of elements of $\mathcal{X}^{V_n}$ with law constructed as follows. First, choose $\boldsymbol{\alpha}$ from the law $\nu^{\times V_n}$. Then, for each $v \in V_n$ independently, choose two random bits $\eta_v, \omega_v \in \{0,1\}$ independently, each equal to 1 with probability κ^s. Finally, for each $v \in V_n$, set

$$\boldsymbol{\xi}_v := \begin{cases} \boldsymbol{\alpha}_v & \text{if } \eta_v = 0 \\ x_v & \text{if } \eta_v = 1 \end{cases} \quad \text{and} \quad \boldsymbol{\zeta}_v := \begin{cases} \boldsymbol{\alpha}_v & \text{if } \omega_v = 0 \\ y_v & \text{if } \omega_v = 1. \end{cases}$$

Letting λ be the joint distribution of $(\boldsymbol{\alpha}, \boldsymbol{\xi}, \boldsymbol{\zeta})$, it follows that

$$\lambda\{\boldsymbol{\xi} \in \cdot\} = \mathbb{P}_n^{\mathbf{x}}\{\boldsymbol{\xi}_{n,s} \in \cdot\}, \quad \lambda\{\boldsymbol{\zeta} \in \cdot\} = \mathbb{P}_n^{\mathbf{y}}\{\boldsymbol{\zeta}_{n,s} \in \cdot\},$$

and

$$\int d^{(V_n)}(\boldsymbol{\zeta}, \boldsymbol{\xi})\, \mathrm{d}\lambda \leqslant 2\kappa^s < \delta/2.$$

Thus, under λ, the pair of random variables $(\boldsymbol{\xi}, \boldsymbol{\zeta})$ are a coupling of $(\boldsymbol{\xi}_{n,s}, \boldsymbol{\zeta}_{n,s})$ under which the probability of the event $\{d^{(V_n)}(\boldsymbol{\xi}_{n,x}, \boldsymbol{\zeta}_{n,s}) < \delta\}$ is greater that $1/2$, by Chebyshev's Inequality. Now we can choose any extension of this to a coupling $\mathbb{Q}$ of the whole random trajectories $(\boldsymbol{\xi}_{n,0}, \dots, \boldsymbol{\xi}_{n,s})$ and $(\boldsymbol{\zeta}_{n,0}, \dots, \boldsymbol{\zeta}_{n,s})$: for instance, we can couple them relatively independently over the given coupling of the end-states $\boldsymbol{\xi}_{n,s}$ and $\boldsymbol{\zeta}_{n,s}$. □

Remark. Using Chernoff's Inequality for the random sum $\sum_v(\eta_v + \omega_v)$, one can actually improve (14) to a lower bound of the form $1 - O(\mathrm{e}^{-\beta|V_n|})$, but we will not need this. $\lhd$

Proof of Proposition 5.2. Let $\mathcal{O}$ be a hereditary w*-neighbourhood of $\nu^{\times G}$ and let $\delta \in (0,1)$. Choose some $\kappa \in (1-\delta, 1)$, and then choose $s \in \mathbb{N}$ so that $\kappa^s < \delta/4$.

Having chosen s, let n be so large that

$$\inf_{\mathbf{x}\in\Omega(\mathcal{O},\sigma_n)} \mathbb{P}_n^{\mathbf{x}}\big\{\boldsymbol{\xi}_{n,t} \in \Omega(\mathcal{O},\sigma_n)\ \forall t = 1,2,\ldots,s\big\} > 3/4 \tag{15}$$

and

$$\inf_{\mathbf{x}\in\mathcal{X}^{V_n}} \mathbb{P}_n^{\mathbf{x}}\big\{d^{(V_n)}(\boldsymbol{\xi}_{n,t}, \boldsymbol{\xi}_{n,t+1}) < \delta\ \forall t = 0,1,\ldots,s-1\big\} > 3/4. \tag{16}$$

This is possible by Lemmas 5.3 and 5.4, respectively. We will show that $\Omega(\mathcal{O},\sigma_n)$ is δ-connected for any such n.

Suppose that $\mathbf{x}, \mathbf{y} \in \Omega(\mathcal{O},\sigma_n)$, and let $\mathbb{Q}$ be a coupling of two trajectories of the random walk up to time s, one starting from $\mathbf{x}$ and the other from $\mathbf{y}$, as constructed in Lemma 5.5. Then the conjunction of the bounds (15), (16) and (14) shows that the event

$$\Big\{\mathbf{x} = \boldsymbol{\xi}_{n,0}, \boldsymbol{\xi}_{n,1}, \ldots, \boldsymbol{\xi}_{n,s-1}, \boldsymbol{\xi}_{n,s}, \boldsymbol{\zeta}_{n,s}, \boldsymbol{\zeta}_{n,s-1}, \ldots, \boldsymbol{\zeta}_{n,1}, \boldsymbol{\zeta}_{n,0} = \mathbf{y} \text{ is a } \delta\text{-path which stays within } \Omega(\mathcal{O},\sigma_n)\Big\}$$

has positive probability under $\mathbb{Q}$: so, in particular, such a δ-path exists. Since δ was arbitrary, this completes the proof. $\square$

Remark. Since the choice of s in the above proof depends only on δ, it actually shows that Bernoulli shifts have uniformly connected model spaces rel Σ, as in the remark at the end of Subsection 2.2. $\lhd$

6 Analysis of Popa factors

6.1 Actions and cocycles for property-(T) groups

Let G be a finitely generated group. Let S be a finite and symmetric generating set, and let R be the set of all the corresponding relations in the free group on S (including concatenations or conjugates of other relations). Recall that G has

Kazhdan's **property (T)** if there is a $c > 0$ for which the following holds: whenever $\pi : G \curvearrowright V$ is a unitary representation, if there is some $v \in V$ such that $\|v\| = 1$ and

$$\max_{s \in S} \|\pi^s v - v\| \leqslant c, \tag{17}$$

then π has a nontrivial invariant vector. The value of c depends on the choice of S, but its existence does not. See, for instance, [2].

Now suppose that (X, μ, T) is a G-system and that $K \leqslant \mathbb{T}$ is a closed subgroup. Let $\mathcal{U}(\mu, K)$ be the set of measurable functions $X \longrightarrow K$ modulo agreement μ-a.e. This is a group under pointwise addition, and is naturally equipped with the topology of convergence in probability. That topology is Polish, with a suitable metric given by the group-norm

$$\|f\|_\mu := \int |f| \, \mathrm{d}\mu,$$

where $|\cdot|$ is the quotient group-norm on $\mathbb{T}$ of the usual absolute value on $\mathbb{R}$. The action $T : G \curvearrowright (X, \mu)$ induces an action of G on $\mathcal{U}(\mu, K)$. We will need to work with the cohomology of this action in degree 1, which is conveniently expressed in terms of the generators S and relations R.

Firstly, a K**-valued** 1**-cochain** is an equivalence class modulo μ of measurable functions $\alpha : S \times X \longrightarrow K$, or equivalently an element of $\mathcal{U}(\mu, K)^S$. For a 1-cochain α, we set

$$\|\alpha\|_{\mu,S} := \sum_{s \in S} \|\alpha(s, \cdot)\|_\mu.$$

This defines a Polish group-norm on the group of 1-cochains under pointwise addition, similarly to $\mathcal{U}(\mu, K)$.

Next, given a 1-cochain $\alpha : S \times X \longrightarrow K$ and a word $w = s_\ell s_{\ell-1} \dots s_1$ over the alphabet S, define

$$\alpha(w, \cdot) := \alpha(s_1, \cdot) + \alpha(s_2, T^{s_1}(\cdot)) + \cdots + \alpha(s_\ell, T^{s_1 \dots s_{\ell-1}}(\cdot)).$$

Restricting to relations, the resulting function

$$R \times X \longrightarrow K : (w, x) \mapsto \alpha(w, x)$$

is the 2**-coboundary** of α, denoted by $d\alpha$. For any $\gamma : R \times X \longrightarrow \mathbb{T}$ and finite $F \subseteq R$, we set

$$\|\gamma\|_{\mu,F} := \max_{w \in F} \|\gamma(w, \cdot)\|_\mu.$$

A 1-cochain is a 1-**cocycle** if $d\alpha = 0$ a.s. These form a closed subgroup $Z^1(T, \mu, K)$ of the group of 1-cochains.

Lastly, if $\beta \in \mathcal{U}(\mu, K)$, then its 1-**coboundary** is the 1-cocycle

$$d\beta : (s, x) \mapsto \beta(T^s x) - \beta(x).$$

These form a further subgroup $B^1(T, \mu, K) \leqslant Z^1(T, \mu, K)$, not necessarily closed, and the quotient of these groups is the **first cohomology group** $H^1(T, \mu, K)$.

The next result is due to Schmidt [20, Theorem 3.4] and independently to Zimmer [23, Theorem 2.11]. We include a proof in order to show that the relevant constants do not depend on the action T, or on the choice of the closed subgroup K of $\mathbb{T}$.

Theorem 6.1. *If G has property (T), then there is some $r > 0$ with the following property. Let (X, μ, T) be an ergodic G-system. For any closed subgroup $K \leqslant \mathbb{T}$ and any $\alpha \in Z^1(T, \mu, K)$, we have*

$$\|\alpha\|_{\mu,S} \leqslant r \quad \Longrightarrow \quad \alpha \in B^1(T, \mu, K).$$

Proof. Step 1. First suppose that $K = \mathbb{T}$.

Let c be the constant in the definition of property (T), let $r := c^2/4\pi$, and consider $\alpha \in Z^1(T, \mu, \mathbb{T})$ with $\|\alpha\|_{\mu,S} \leqslant r$. Let $\pi : G \curvearrowright L^2(\mu)$ be the Koopman representation of T twisted by α:

$$(\pi^g f)(x) := \mathrm{e}^{2\pi \mathrm{i} \alpha(g^{-1}, x)} \cdot f(T^{g^{-1}} x).$$

This is a well-defined G-action because α satisfies the defining equations for a 1-cocycle.

Now observe that

$$\|\pi^s 1_X - 1_X\|_{L^2(\mu)} = \sqrt{\int |\mathrm{e}^{2\pi \mathrm{i} \alpha(s^{-1}, x)} - 1|^2 \, \mu(\mathrm{d}x)} \leqslant \sqrt{4\pi r} = c,$$

where 1_X is the constant function 1 on X. Therefore property (T) gives some $f \in L^2(\mu)$ such that $\|f\|_{L^2(\mu)} = 1$ and $\pi^g f = f$ for all $g \in G$.

This fixed-point equation implies that $|f|$ is T-invariant, hence μ-a.s. constant by ergodicity. Therefore f is actually S^1-valued, and the fixed-point equation reads

$$\mathrm{e}^{2\pi \mathrm{i} \alpha(g, x)} = \overline{f(T^g x)} f(x) \quad \forall g \in G.$$

Taking arguments, this asserts that $\alpha \in B^1(T, \mu, \mathbb{T})$, as required.

Step 2. Now consider general $K \leqslant \mathbb{T}$. If we regard α as a $\mathbb{T}$-valued function, Step 1 gives some $\beta_0 \in \mathcal{U}(\mu, \mathbb{T})$ such that

$$\alpha(g, x) = \beta_0(T^g x) - \beta_0(x) \quad \text{for } \mu\text{-a.e. } x \; \forall g \in G.$$

Since α is actually K-valued, this implies that the coset $\beta_0(x) + K$ is T-invariant, and hence a.s. constant by ergodicity. Let $c + K$ be that coset, and let $\beta(x) := \beta_0(x) - c$. Then β takes values in K almost surely, and still satisfies

$$\alpha(g, x) = \beta(T^g x) - \beta(x) \quad \text{for } \mu\text{-a.e. } x \; \forall g \in G.$$

□

Next we introduce a 'roughened' notion of cocycles. Given a G-system (X, μ, T), $\varepsilon > 0$, and finite $F \subseteq R$, an element $\alpha \in \mathcal{U}(\mu, K)^S$ is a K**-valued** (F, ε)**-near-cocycle over** (X, μ, T) if

$$\|d\alpha\|_{\mu,F} < \varepsilon.$$

The set of these will be denoted $Z^1_{F,\varepsilon}(T, \mu, K)$. For these we have the following roughened version of Theorem 6.1.

Theorem 6.2. *Suppose $G = \langle S \mid R \rangle$ has property (T), and let $r > 0$ be as Theorem 6.1. Then for every $r' > 0$ there are $\varepsilon > 0$ and finite $F \subseteq R$ such that the following holds. If (X, μ, T) is an ergodic G-system, and $\alpha \in Z^1_{F,\varepsilon}(T, \mu, K)$ satisfies*

$$\|\alpha\|_{\mu,S} \leqslant r,$$

then there is some $\beta \in \mathcal{U}(\mu, K)$ such that

$$\|\alpha - d\beta\|_{\mu,S} < r'.$$

Proof. Step 1. We first convert the result to an assertion about invariant measures on a fixed G-space. Let $Y := (K^S)^G = K^{S \times G}$, equipped with the G-action by coordinate right-shift, and define the canonical 1-cochain $\alpha_0 : S \times Y \longrightarrow K$ by

$$\alpha_0(s, y) := y_{s,e}.$$

We will prove that the desired result is implied by the following:

> *For every $r' > 0$ there are $\varepsilon > 0$ and finite $F \subseteq G$ such that the following holds. If ν is an ergodic shift-invariant Borel probability on Y, and if*
>
> $$\|\alpha_0\|_{\nu,S} \leqslant r \quad \textit{and} \quad \|d\alpha_0\|_{\nu,F} < \varepsilon,$$

then there is some $\beta_0 \in \mathcal{U}(\nu, K)$ *such that*

$$\|\alpha_0 - d\beta_0\|_{\nu,S} < r'.$$

Indeed, suppose this result is known, and consider (X, μ, T) and α as in the statement of the theorem. Define

$$\varphi : X \longrightarrow Y : x \mapsto (\alpha(s, T^g x))_{s\in S, g\in G}.$$

Then φ intertwines T with the coordinate shift on Y, and so $\nu := \varphi_*\mu$ is an ergodic G-invariant probability on Y. Moreover, the definition of φ gives $\alpha = \alpha_0 \circ \varphi$, and so

$$\|\alpha_0\|_{\nu,S} = \|\alpha\|_{\mu,S} \leqslant r \quad \text{and} \quad \|d\alpha_0\|_{\nu,F} = \|d\alpha\|_{\mu,F} < \varepsilon.$$

So the claim above provides some measurable $\beta_0 : Y \longrightarrow K$ such that, setting $\beta := \beta_0 \circ \varphi$, we have

$$\|\alpha - d\beta\|_{\mu,S} = \|\alpha_0 - d\beta_0\|_{\nu,S} < r'.$$

Step 2. The rest of the proof is by contradiction. Fix an increasing sequence $(F_i)_{i\geqslant 1}$ of finite sets whose union is R, and suppose that one could find $r' > 0$ and a sequence $(\nu_i)_{i\geqslant 1}$ of ergodic shift-invariant Borel probabilities on Y such that

$$\|\alpha_0\|_{\nu_i,S} \leqslant r \quad \text{and} \quad \|d\alpha_0\|_{\nu_i,F_i} < 2^{-i} \tag{18}$$

for all i, but also such that

$$\|\alpha_0 - d\beta_0\|_{\nu_i,S} \geqslant r'. \tag{19}$$

for all measurable functions $\beta_0 : Y \longrightarrow K$ and all i.

By passing to a subsequence, we may assume that these measures ν_i weak*-converge to a measure ν, which must still be shift-invariant. Since G has property (T), a theorem of Glasner and Weiss [7] gives that the set of ergodic measures is weak*-closed among the shift-invariant probability measures on Y, so this ν is still ergodic.

Now, since each $\alpha_0(s, \cdot)$ is continuous on Y, the two parts of (18) give

$$\|\alpha_0\|_{\nu,S} = \lim_{i\longrightarrow\infty} \|\alpha_0\|_{\nu_i,S} \leqslant r$$

and

$$\|d\alpha_0(w, \cdot)\|_\nu = \lim_{i\longrightarrow\infty} \|d\alpha_0(w, \cdot)\|_{\nu_i} = 0 \quad \forall w \in \bigcup_i F_i = R.$$

Therefore $\alpha_0 : S \times Y \longrightarrow K$ is an element of $Z^1(\mathrm{shift}, \nu, K)$ with $\|\alpha_0\|_{\nu,S} \leqslant r$, so Theorem 6.1 gives some $\beta_1 \in \mathcal{U}(\nu, K)$ such that

$$\alpha_0 = d\beta_1 \quad \nu\text{-a.s..}$$

Letting β_0 be a continuous function that approximates β_1 sufficiently well in ν-probability, we can arrange that

$$\|\alpha_0 - d\beta_0\|_{\nu,S} < r'.$$

Since β_0 is continuous and $\nu_i \overset{\text{weak}^*}{\longrightarrow} \nu$, this implies that also

$$\|\alpha_0 - d\beta_0\|_{\nu_i,S} < r'.$$

for all sufficiently large i. This contradicts (19). □

Since the bounds in Theorem 6.2 do not depend on the system (X, μ, T), they are nontrivial even for actions on finite sets. In particular, suppose that $H < G$ is a finite-index subgroup, let $V := G/H$, and let $\Gamma := (V, E)$ be the directed Schreier graph resulting from the generating set S. Let G act on V by left-multiplication. Observe that $S \times V$ is in canonical bijection with E, so a 1-cochain for this system may be interpreted as a map $\alpha : E \longrightarrow K$. For any $F \subseteq R$, let L_F be the set of based loops in Γ that correspond to walking around a relation from the set F, starting from any vertex. We will identify such a loop by a pair (v, w), where v is its starting vertex and w is the relation. For any $\alpha : E \longrightarrow K$, we can now define $d\alpha : L_R \longrightarrow K$ by

$$d\alpha(v, w) := \sum_{i=0}^{\ell-1} \alpha(s_i \cdots s_1 v, s_{i+1} s_i \cdots s_1 v),$$

where $w = s_\ell s_{\ell-1} \cdots s_1$. As before, a map $\alpha : E \longrightarrow K$ is a 1-cocycle if and only if $d\alpha = 0$: that is, it sums to zero around any loop in Γ which corresponds to a relation of the group.

Similarly to the case of general G-systems, let $Z^1(\Gamma, K)$ be the set of 1-cocycles $\alpha : E \longrightarrow K$, let $B^1(\Gamma, K) = d(K^V)$ be the subgroup of coboundaries, and for $\varepsilon > 0$ and $F \subseteq R$ let $Z^1_{\varepsilon,F}(\Gamma, K)$ be the subset of 1-cochains α which satisfy

$$\frac{1}{|V|} \sum_{v \in V} |d\alpha(v, w)| < \varepsilon \quad \forall w \in F.$$

Let d be the metric on $\mathbb{T}^S$ given by

$$d(\theta_1, \theta_2) := \sum_{s \in S} |\theta_{1,s} - \theta_{2,s}|.$$

We also write d for the restriction of this metric to K^S for any closed subgroup K of $\mathbb{T}$. Since $S \times V$ is canonically identified with E, the normalized Hamming sums $d^{(V)}$ may be regarded as metrics on $\mathbb{T}^E$. Now the translation of Theorems 6.1 and 6.2 into this setting implies the following.

Corollary 6.3. *Let r be as in Theorem 6.1.*

1. *If $\alpha \in Z^1(\Gamma, K)$ and $d^{(V)}(\alpha, B^1(\Gamma, K)) \leqslant r$ then in fact $\alpha \in B^1(\Gamma, K)$.*
2. *Given $r' > 0$, let $\varepsilon > 0$ and $F \subseteq R$ be as in Theorem 6.2. If $\alpha \in Z^1_{\varepsilon,F}(\Gamma, K)$ and $d^{(V)}(\alpha, B^1(\Gamma, K)) \leqslant r$ then in fact*

$$d^{(V)}(\alpha, B^1(\Gamma, K)) < r'.$$

Proof. Suppose $d\beta \in B^1(\Gamma, K)$ is such that $d^{(V)}(\alpha, d\beta) \leqslant r$. It suffices to prove either part with α replaced by $\alpha - d\beta$, and so we may simplify the assumption to $d^{(V)}(\alpha, 0) \leqslant r$. What remains is a special case of Theorem 6.2. □

Remark. The relation between ε and r' in part 2 of this corollary is reminiscent of recent work on coboundary expansion for simplicial complexes by Kaufman, Kazhdan and Lubotzky [10]. They prove a similar inequality for $K = \mathbb{Z}/2\mathbb{Z}$ and for a family of graphs Γ constructed from certain Bruhat-Tits buildings. Their family of examples is much more specialized than ours, but they obtain a more quantitative result: a linear dependence of ε on r'. They need this extra precision to deduce some other expansion properties of these complexes. It would be interesting to know whether one can be so precise under our more general conditions. ◁

6.2 Proof of Theorem D

Now let G be an infinite, finitely generated, residually finite group having Kazhdan's property (T), and let e be its identity element. Let S and R be as in the previous subsection. Such a group G has a descending sequence $G_1 > G_2 > \dots$ of finite-index subgroups such that $\bigcap_n G_n = \{e\}$.

For the proof of Theorem D we will need those subgroups to have two additional properties.

Lemma 6.4. *There is a descending sequence $(G_n)_{n\geqslant 1}$ of finite-index subgroups as above such that*

(i) they all have nontrivial Abelianization, and

(ii) the left-multiplication actions $\sigma_n : G \longrightarrow \mathrm{Sym}(G/G_n)$ form a sofic approximation to G.

Proof. First, as is standard, there is a sequence $G > H_1 > H_2 > \dots$ of finite-index *normal* subgroups converging to $\{e\}$.

For each n, consider the quotient homomorphism

$$q_n : H_n \twoheadrightarrow H_n/H_{n+1}.$$

Since the target is a nontrivial group, it contains a nontrivial cyclic subgroup $K_n \leqslant H_n/H_{n+1}$. Letting $G_n := q_n^{-1}(K_n)$, we obtain $H_n \geqslant G_n > H_{n+1}$. Let σ_n be the left-multiplication action of G on G/G_n.

We can now deduce the required properties.

(i) By construction, K_n is a quotient group of $G_n/[G_n, G_n]$, so this latter is nontrivial for every n.

(ii) If left-multiplication by $g \in G$ has a fixed point in G/G_n, then it has a fixed point in the further quotient G/H_n. Since H_n is normal in G, this requires that $g \in H_n$. Since the subgroups H_n converge to $\{e\}$, it follows that, for any $g \neq e$, the permutation σ_n^g has no fixed points for all sufficiently large n. This implies that $(\sigma_n)_{n\geqslant 1}$ is a sofic approximation.

□

Henceforth G, S and the sequence $(G_n)_n$ given by the above lemma will be fixed. Let $\Gamma_n := (V_n, E_n)$ be the sequence of directed Schreier graphs on the vertex sets $V_n = G/G_n$ and with E_n defined by the generating set S. These form a sofic approximation to G via the homomorphisms σ_n in part (ii) of the above lemma. In this setting we often write $g \cdot v$ in place of $\sigma_n^g(v)$ for $g \in G$ and $v \in V_n$.

Now let us return to the Popa factor map described in the Introduction. An isomorphic factor map whose target is a shift-system may be constructed as follows:

$$\Phi : \mathbb{T}^G \longrightarrow \mathbb{T}^{S\times G} : (\theta_g)_{g\in G} \mapsto (\theta_{sg} - \theta_g)_{s\in S, g\in G}.$$

This equals φ^G for the map

$$\varphi : \mathbb{T}^G \longrightarrow \mathbb{T}^S : \theta \mapsto (\theta_s - \theta_e)_{s\in S}.$$

Since φ depends on only finitely many coordinates in $\mathbb{T}^G$, it is an η-AL approximation to itself for every $\eta > 0$, and we may work directly with the maps φ^{σ_n} on model spaces.

This Φ is a homomorphism of compact Abelian groups. Since S generates G, the kernel of Φ is the diagonal subgroup of $\mathbb{T}^G$. The image of Φ is a compact subgroup of $\mathbb{T}^{S\times G}$, which we denote by Z. The image measure $\Phi_* m^{\times G}$ must equal the Haar measure μ_Z of Z. Therefore Φ is a factor map

$$(\mathbb{T}^G, m^{\times G}, S) \longrightarrow (\mathbb{T}^{S\times G}, \mu_Z, S)$$

which is equivalent to the Popa factor map from the Introduction.

To prove Theorem D, we will need a description of the spaces of good models for this factor system. This begins with the following, which is an immediate consequence of the definitions.

Lemma 6.5. *For any finite $F \subseteq R$ and $\varepsilon > 0$ there is a w*-neighbourhood $\mathcal{O}$ of μ_Z such that*

$$\Omega(\mathcal{O}, \sigma_n) \subseteq Z^1_{\varepsilon,F}(\Gamma_n, \mathbb{T}) \quad \forall n \geqslant 1.$$

□

If $E \subseteq G$ and $\pi_E : \mathbb{T}^{S\times G} \longrightarrow \mathbb{T}^{S\times E}$ is the coordinate projection, then π_E is a group homomorphism, and the image measure $(\mu_Z)_E = (\pi_E)_* \mu_Z$ equals the Haar measure $\mu_{\pi_E(Z)}$.

Lemma 6.6. *Let $E \subseteq G$ be finite, let $C \subseteq G$ be a finite subset which is connected in $\mathrm{Cay}(G, S)$ and such that $C \supseteq E \cup SE$, and let n be so large that, for any $v \in V_n$, the map*

$$C \longrightarrow V_n : g \mapsto g \cdot v$$

is a graph isomorphism between the restriction of $\mathrm{Cay}(G, S)$ to C and the restriction of Γ_n to $C \cdot v$. Then for any $\alpha \in Z^1(\Gamma_n, \mathbb{T})$ and any $v \in V_n$ we have

$$\left\{(\gamma(s, g \cdot v))_{s\in S, g\in E} : \ \gamma \in \alpha + B^1(\Gamma_n, \mathbb{T})\right\} = \pi_E(Z). \tag{20}$$

Proof. We first reduce to the case $\alpha = 0$. To this end, observe that if $\alpha : E_n \longrightarrow \mathbb{T}$ is a cocycle, then it has zero sum around any based loop in Γ_n which corresponds to a relation of the group. The map

$$C \longrightarrow C \cdot v : g \mapsto g \cdot v \tag{21}$$

defines an isomorphism of the restricted graphs by assumption, so any loop of Γ_n that is contained in $C \cdot v$ must arise from a relation of the group. Therefore the cocycle condition implies that α sums to zero around any loop contained in $C \cdot v$, and hence

$$(\alpha|_{S\times g\cdot v})_{g\in E} = ((d\beta)|_{S\times g\cdot v})_{g\in E}$$

for some $\beta \in \mathbb{T}^{C\cdot v}$ obtained by simply summing along paths from some distinguished basepoint in C: all of C can be reached this way because C is connected. Extending β arbitrarily to a member of $\mathbb{T}^{V_n}$, this shows that the left-hand side of (20) does not depend on α: for every α that left-hand side is a coset of a homomorphic image of $B^1(\Gamma_n, \mathbb{T})$, and we have just seen that any two of these cosets overlap, hence are equal.

Now suppose that $\alpha = 0$. We will prove two separate inclusions.

First, if $\theta \in \mathbb{T}^G$ and

$$\gamma = \pi_E(\Phi(\theta)) = (\theta_{sg} - \theta_g)_{s\in S, g\in E} \in \pi_E(Z),$$

then γ is also equal to $((d\beta)|_{S\times g\cdot v})_{g\in E}$ for any choice of $\beta \in \mathbb{T}^{V_n}$ satisfying

$$\beta_{g\cdot v} = \theta_g \quad \forall g \in E.$$

Such a choice exists because $C \supseteq E$ and the map (21) is injective.

On the other hand, if $\beta \in \mathbb{T}^{V_n}$, then the reverse of this argument produces some $\theta \in \mathbb{T}^G$ such that $\beta_{g\cdot v} = \theta_g$ for all $g \in E$. □

Lemma 6.7. *For any w*-neighbourhood $\mathcal{O}$ of μ_Z, we have*

$$\inf_{\alpha\in Z^1(\Gamma_n,\mathbb{T})} ((\varphi^{\sigma_n})_* m^{\times V_n})\big(\Omega(\mathcal{O}, \sigma_n) - \alpha\big) \longrightarrow 1$$

as $n \longrightarrow \infty$.

Proof. For $\alpha \in Z^1(\Gamma_n, \mathbb{T})$, let

$$R_\alpha : \mathbb{T}^{V_n} \longrightarrow \mathbb{T}^{V_n} : \beta \mapsto \beta + \alpha$$

be the corresponding rotation. We need to show that the image measures

$$(R_{\alpha_n})_*(\varphi^{\sigma_n})_* m^{\times V_n}$$

are asymptotically supported on $\Omega(\mathcal{O}, \sigma_n)$ as $n \longrightarrow \infty$ for any sequence $\alpha_n \in Z^1(\Gamma_n, \mathbb{T})$. We will deduce this by showing that

$$(R_{\alpha_n})_*(\varphi^{\sigma_n})_* m^{\times V_n} \xrightarrow{\mathrm{q}} \mu_Z,$$

where this refers to quenched convergence as in [1, Definition 5.3].

Since $(\mathbb{T}^{S\times G}, \mu_Z, S)$ is a factor of a Bernoulli shift, it is ergodic. Therefore by [1, Corollary 5.7] quenched convergence will follow if we show that

$$(R_{\alpha_n})_*(\varphi^{\sigma_n})_* m^{\times V_n} \xrightarrow{\text{lw}*} \mu_Z \quad \text{(local weak* convergence)}.$$

The measure $(R_{\alpha_n})_*(\varphi^{\sigma_n})_* m^{\times V_n}$ is equal to the Haar measure on the coset

$$B^1(\Gamma_n, \mathbb{T}) + \alpha_n.$$

The result now follows from Lemma 6.6. According to that lemma, for any finite $E \subseteq G$, the projection of this coset to the directed edges which emanate from $E \cdot v$ is simply a copy of $\pi_E(Z)$, provided n is large enough depending on E. For such n, the projection of $(R_{\alpha_n})_*(\varphi^{\sigma_n})_* m^{\times V_n}$ to $E \cdot v$ is therefore equal to $\mu_{\pi_E(Z)} = (\mu_Z)_E$. □

Proof of Theorem D. Let r be as in Theorem 6.1, let $\delta := r/4$, and let $\varepsilon > 0$ and $F \subseteq R$ be given by part 2 of Corollary 6.3 for $r' := r/10$. Finally, let $\mathcal{O}$ be a w*-neighbourhood of μ_Z as given by Lemma 6.5 for this F and ε.

For any w*-neighbourhood $\mathcal{U}$ of μ_Z, Lemma 6.7 gives that

$$\Omega(\mathcal{U}, \sigma_n) \cap (\alpha + B^1(\Gamma_n, \mathbb{T})) \neq \varnothing \tag{22}$$

for any $\alpha \in Z^1(\Gamma_n, \mathbb{T})$, once n is sufficiently large.

Next, for each n, we may identify $G \curvearrowright \mathbb{T}^{G/G_n}$ as the induction to G of the trivial action of G_n on $\mathbb{T}$. Hence Shapiro's Lemma [15, Theorem II.3.7] gives

$$H^1(G, \mathbb{T}^{G/G_n}) \cong H^1(G_n, \mathbb{T}) = \mathrm{Hom}(G_n, \mathbb{T}) = \big(G_n/[G_n, G_n]\big)^\wedge \neq 0,$$

where the last conclusion is the point at which we use part (i) of Lemma 6.4. Therefore for each n there is more than one coset of $B^1(\Gamma_n, \mathbb{T})$ in $Z^1(\Gamma_n, \mathbb{T})$.

Combining this with (22), it follows that we may choose sequences

$$\alpha_n \in B^1(\Gamma_n, \mathbb{T}) \quad \text{and} \quad \alpha_n' \in Z^1(\Gamma_n, \mathbb{T}) \backslash B^1(\Gamma_n, \mathbb{T})$$

satisfying

$$P^{\sigma_n}_{\alpha_n}, P^{\sigma_n}_{\alpha_n'} \xrightarrow{\text{weak}^*} \mu_Z \quad \text{as } n \longrightarrow \infty.$$

By part 1 of Corollary 6.3, we must have that

$$d^{(V_n)}(\alpha_n', B^1(\Gamma_n, \mathbb{T})) \geqslant r$$

for all sufficiently large n.

Now suppose, for the sake of contradiction, that there were δ-paths

$$\alpha_n = \alpha_{n,1}, \alpha_{n,2}, \ldots, \alpha_{n,\ell_n} = \alpha'_n$$

contained in $\Omega(\mathcal{O}, \sigma_n)$ for arbitrarily large n. By Lemma 6.5, they would also be contained in $Z^1_{\varepsilon,F}(\Gamma_n, \mathbb{T})$. Since $\delta = r/4$, there would have to be some $k \in \{2, \ldots, \ell_n - 1\}$ for which

$$r' < r/4 < d^{(V_n)}(\alpha_{n,k}, B^1(\Gamma_n, \mathbb{T})) < 3r/4.$$

Since $\alpha_{n,k} \in Z^1_{\varepsilon,F}(\Gamma_n, \mathbb{T})$, this would contradict part 2 of Corollary 6.3. □

6.3 Complemented factors

The deduction of Corollary D″ rests on the following theorem about a general metric G-process $(\mathcal{X}^G, \mu, S, d_\mathcal{X})$, which may be of independent interest. It is expressed in terms of a sequence of measures on the spaces $\mathcal{X}^{V_n}$ which 'doubly-quenched converge' to $\mu \in \mathrm{Prob}^S(\mathcal{X}^G)$ over Σ, denoted by $\mu_n \xrightarrow{\mathrm{dq}} \mu$. This property is introduced in [1, Subsection 5.2]. By [1, Corollary 5.18], the existence of such a sequence μ_n is an isomorphism-invariant of the process which does not depend on the metric $d_\mathcal{X}$. We do not recall the definition here, but refer the reader to that reference for full details.

Theorem 6.8. *Suppose there exist $\mu_n \in \mathrm{Prob}(\mathcal{X}^{V_n})$ such that $\mu_n \xrightarrow{\mathrm{dq}} \mu$ rel Σ. Then any complemented factor map of $(\mathcal{X}^G, \mu, S)$ is model-surjective rel Σ.*

In particular, any complemented factor map of a Bernoulli shift is model-surjective relative to any sofic approximation of G.

Proof. Suppose that $\Phi : (\mathcal{X}^G, \mu, S, d_\mathcal{X}) \longrightarrow (\mathcal{Y}^G, \nu, S, d_\mathcal{Y})$ is a complemented factor map, so there is another factor map $\Psi : (\mathcal{X}^G, \mu, S, d_\mathcal{X}) \longrightarrow (\mathcal{Z}^G, \theta, S, d_\mathcal{Z})$ such that the combined map

$$(\Phi, \Psi) : (\mathcal{X}^G, \mu, S) \longrightarrow (\mathcal{Y}^G \times \mathcal{Z}^G, \nu \times \theta, S)$$

is a measure-theoretic isomorphism.

By [1, Proposition 5.16], there are also measures $\theta_n \in \mathrm{Prob}(\mathcal{Z}^{V_n})$ such that $\theta_n \xrightarrow{\mathrm{dq}} \theta$ rel Σ: they can be obtained as images of the measures μ_n. Using these, [1, Corollary 5.13] gives that the coordinate-projection factor map

$$\Pi : (\mathcal{Y}^G \times \mathcal{Z}^G, \nu \times \theta, S) \longrightarrow (\mathcal{Y}^G, \nu, S)$$

is model-surjective rel Σ. This conclusion is written in terms of w*-neighbourhoods in [1], but it is easily turned into condition 2 of Lemma 3.4. Therefore the factor map

$$\Phi = \Pi \circ (\Phi, \Psi)$$

is a composition of a factor map which is model-surjective rel Σ and an isomorphism, so Φ itself is also model-surjective rel Σ by Propositions 3.5 and 3.6.

This argument applies whenever $(\mathcal{X}^G, \mu, S)$ is a Bernoulli shift and Σ is a sofic approximation, because then the corresponding product measures on the spaces $\mathcal{X}^{V_n}$ always doubly-quenched converge to μ [1, Lemma 5.11]. □

Proof of Corollary D″. This follows at once from the combination of Theorems A, C, D and 6.8. □

Remark. It is easy to prove that the Popa factor map itself is not complemented for any infinite discrete group G. This is because it is a relatively compact extension, so if it had a complementing factor $\Xi : (\mathbb{T}^G, m^{\times G}, S) \longrightarrow (Z, \theta, R)$ then (Z, θ, R) would have to be a compact system, but the original Bernoulli shift $(\mathbb{T}^G, m^{\times G}, S)$ has no compact factors.

Corollary D″ is much stronger: it prohibits the Popa factor system from appearing as a complemented factor of any other Bernoulli system in any way. This holds only for certain special groups G. By contrast, for amenable G the Popa factor is always isomorphic to another copy of the Bernoulli shift by the general theory of [17], and the same holds if G is a free group by an easy calculation. ◁

Acknowledgements

This research was supported partly by a fellowship from the Clay Mathematics Institute and partly by the Simons Collaboration on Algorithms and Geometry

I am grateful for several helpful discussions with Lewis Bowen, David Fisher, Alex Lubotzky and Brandon Seward, and to Brandon Seward and the anonymous referee for correcting several errors in an earlier version. ◁

References

[1] T. Austin. Additivity properties of sofic entropy and measures on model spaces. Preprint, available online at `arXiv.org`: 1510.02392.

[2] B. Bekka, P. de la Harpe, and A. Valette. *Kazhdan's property (T)*, volume 11 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2008.

[3] L. Bowen. Measure conjugacy invariants for actions of countable sofic groups. *J. Amer. Math. Soc.*, 23(1):217–245, 2010.

[4] L. Bowen. Weak isomorphisms between Bernoulli shifts. *Israel J. Math.*, 183:93–102, 2011.

[5] L. Bowen. Every countably infinite group is almost Ornstein. In *Dynamical systems and group actions*, volume 567 of *Contemp. Math.*, pages 67–78. Amer. Math. Soc., Providence, RI, 2012.

[6] L. Bowen. Sofic entropy and amenable groups. *Ergodic Theory Dynam. Systems*, 32(2):427–466, 2012.

[7] E. Glasner and B. Weiss. Kazhdan's property T and the geometry of the collection of invariant measures. *Geometric and Functional Analysis*, 7:917–935, 1997.

[8] M. Gromov. *Metric Structures for Riemannian and Non-Riemannian Spaces*. Birkhäuser, Boston, 2nd edition, 2001.

[9] Y. Jiang. A remark on T-valued cohomology groups of algebraic group actions. Preprint, available online at `arXiv.org`: 1509.08278.

[10] T. Kaufman, D. Kazhdan, and A. Lubotzky. Ramanujan complexes and bounded degree topological expanders. In *55th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2014*, pages 484–493. IEEE Computer Soc., Los Alamitos, CA, 2014.

[11] D. Kerr. Bernoulli actions of sofic groups have completely positive entropy. Preprint, available online at `http://www.math.tamu.edu/~kerr/bernoullicpe20.pdf`.

[12] D. Kerr and H. Li. Bernoulli actions and infinite entropy. *Groups Geom. Dyn.*, 5(3):663–672, 2011.

[13] D. Kerr and H. Li. Entropy and the variational principle for actions of sofic groups. *Invent. Math.*, 186(3):501–558, 2011.

[14] D. Kerr and H. Li. Soficity, amenability, and dynamical entropy. *Amer. J. Math.*, 135(3):721–761, 2013.

[15] S. Lang. *Topics in cohomology of groups*, volume 1625 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1996. Translated from the 1967 French original by the author, Chapter X based on letters written by John Tate.

[16] M. Ledoux. *The concentration of measure phenomenon*, volume 89 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2001.

[17] D. S. Ornstein and B. Weiss. Entropy and isomorphism theorems for actions of amenable groups. *J. Analyse Math.*, 48:1–141, 1987.

[18] S. Popa. Some computations of 1-cohomology groups and construction of non-orbit-equivalent actions. *J. Inst. Math. Jussieu*, 5(2):309–332, 2006.

[19] S. Popa and R. Sasyk. On the cohomology of Bernoulli actions. *Ergodic Theory Dynam. Systems*, 27(1):241–251, 2007.

[20] K. Schmidt. Amenability, Kazhdan's property T, strong ergodicity and invariant means for ergodic group-actions. *Ergodic Theory Dynamical Systems*, 1(2):223–236, 1981.

[21] B. Seward. Krieger's finite generator theorem for actions of countable groups I. Preprint, available online at `arXiv.org`: 1405.3604.

[22] B. Seward. Krieger's finite generator theorem for actions of countable groups II. Preprint, available online at `arXiv.org`: 1501.03367.

[23] R. J. Zimmer. On the cohomology of ergodic actions of semisimple Lie groups and discrete subgroups. *Amer. J. Math.*, 103(5):937–951, 1981.

Courant Institute, New York University
251 Mercer St, New York, NY 10012, U.S.A.
`tim@cims.nyu.edu`